Tracker 3

Object Tracking and Image Processing Software

written by:

Robert Klimek

Ted Wright

1/5/98

Introduction

Tracker is an object tracking and image processing program designed to help with the analysis of images generated by combustion and fluid physics experiments. Experiments are often recorded on film or video tape for later analysis. Tracker automates the process of examining each frame of the recorded experiment, performing image processing operations to bring out the desired detail, and recording the positions of the objects of interest.

Tracker can load sequences of images from disk files, or acquire images (via a frame grabber) from film transports, video tape, laser disks, or a live camera. Tracker controls the image source to automatically advance to the next frame. A large array of image processing operations can be employed to enhance the detail of the acquired images. An arbitrarily large number of objects can be analyzed simultaneously. Several different tracking algorithms are available, including conventional threshold, manual, and correlation based techniques, and more esoteric procedures such as “snake” tracking and automated recognition of character data in the image. Tracker saves results in a text file that can be imported into other programs for graphing or further analysis.

What’s new in this version

- Tracker 3 takes advantage of newer inexpensive ($500) frame grabbers (as opposed to the $16000 frame grabber required for previous versions of Tracker).
- Tracker 3 is designed to be useful even if it is run on a computer without a frame grabber; images can be copied from tape or film to disk files at a tracking station and analyzed later on desktop computers. Previous versions of Tracker would only run on dedicated tracking stations.

- Tracker 3 uses an image processing library that does not require an license for each copy that is run; previous versions of Tracker required the purchase of a $4000 license for each station.
- Previous versions of Tracker were designed for the DOS or Windows 3 operating systems. Tracker 3 is designed for the Windows 95 and Windows NT operating systems.
- Tracker 3 has a new user interface that is easier to use, and that allows combining image processing and tracking techniques in new ways that should make it possible to track new types of objects.

Getting started

Tracker is currently available on diskettes and in a self-extracting executable file available via FTP. A computer running Windows 95 or Windows NT (versions prior to 4.0 have not been tested) is required. Currently, the only supported frame grabbers are the Matrox Meteor (for Pentium class computers) and the Matrox Meteor/PPB for Pentium Pro class computers. Future support is planned for a higher resolution Matrox frame grabber.

Software Installation

For the diskette distribution of Tracker, run the setup program (“setup.exe”) on the first disk.

For the self-extracting executable distribution (“Tracker3.exe”), run the executable in a temporary directory to extract the setup files, then run the setup program (“setup.exe”). The setup files can then be deleted.

The setup program includes a “custom installation” option that allows the choice of whether or not sample images and source code will be installed. Tracker will work without any of these. A typical installation of Tracker (with sample images, but no source code) requires about 7Meg of disk space.

After setup, there will be a Tracker icon on your desktop that you can double-click to start Tracker. There will also be a Tracker entry on your Windows “Start” button.

If you need to un-install Tracker, this can be done from the “Add/Remove Programs” section of the Windows Control Panel.

Basic concepts

Tracker is designed to track “objects” in video and film images. The “objects” that Tracker looks for are assumed to be fairly large (roughly 10 pixels or preferably more in diameter) compared to the objects tracked by “particle tracking” programs. If you want to track 500 2-pixel dots, you probably should not be using Tracker.

Tracker helps to automate the process of retrieving a series of single frame images from film, laserdisk, or tape. It does this by using a serial interface to control “transport devices”. Tracker does not work with every device. The VCRs (for example) that Tracker works with were designed with computer control in mind, and they provide a serial interface for computer connection (and their manufacturers made the necessary software protocol information available so that Tracker can be programmed to control that specific model of VCR). They also reliably advance by a single frame and have a stable picture while paused. Most VCRs do not meet these qualifications. The degree to which specific transport devices work well with computers is highly variable. Tracking from disk files is the most reliable, followed by laserdisk, the film transports, one specific S-VHS VCR (JVC 525), and lastly all other S-VHS and Hi8 VCRs.

Tracker can automate the copying of images from any of the serially controlled transports to disk files. This is recommended if a lot of analysis is required (as an additional benefit, once your images are on disk you can analyze them on a computer without a frame grabber). Another option is to copy your tape to laserdisk (using the built-in analog play and record capacities of the devices), then use Tracker to track from laserdisk.

Tracker analyzes an image using one or more “areas of interest”, or “AOIs”. An AOI is a subset of the image that has a set of image processing algorithms and a tool or tracking algorithm associated with it. The image processing algorithms run in sequence to enhance the image, then the tool or tracking algorithm is applied. Tool algorithms are usually used interactively to determine the best image processing algorithms for enhancing detail in the images.

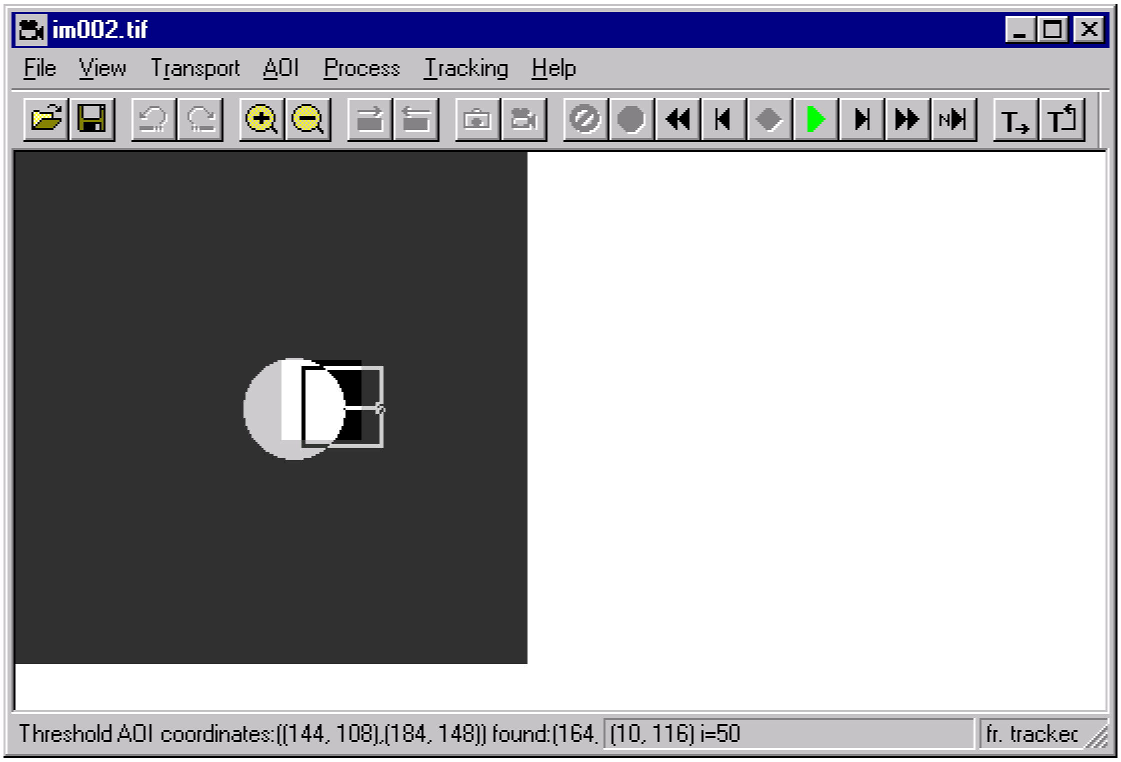

Figure 1. – Simple tracking example.

Simple tracking example

This is a brief example of an easy tracking session to introduce some of the Tracker concepts.

1. Double-click the Tracker icon to start the program.

2. In this example, we will be tracking from the sample images installed with Tracker. To open the first of these images, select "File→Open" from the Tracker menu, double-click the "Images" folder to look in that directory, then double-click the "im001.tif" file to load that file into Tracker.

3. Notice that some of the buttons toward the right side of Tracker have become active. These buttons (and the corresponding items on the "Transport" menu - all buttons have corresponding menu items) are similar to the buttons on a tape deck. There are buttons for advancing frames, "playing" a sequence of frames, "rewinding" to the start of the sequence, etc. Try some of these buttons, and watch how they affect the sequence of images. Also notice the file name of the currently displayed image in the tile of the Tracker window. If you were tracking from tape instead of disk files, these buttons would be controlling the actual tape deck. Rewind back to the first image when you are finished experimenting.

4. Tracker does not normally examine all of an image; it usually works with a only a small portion called the "area of interest" or AOI. Create a new Threshold Tracking AOI by clicking on the "AOI→New→Tracking→Threshold Tracking" menu item. Notice the rectangle with the small arrow inside it that appears in the image. This is a Threshold Tracking AOI. Tracker will only look at image pixels that are inside of this rectangle. Clicking on the corner of the AOI and dragging the mouse will change the size of the AOI, and clicking and dragging on the end of the arrow will change the tracking direction. These points are called the "control points" of the AOI. You don't want to do those things yet, so click and drag the mouse somewhere *not* near the control points to move the whole AOI. This example will track the right edge of the ball in the image. Drag the Threshold Tracking AOI so that it is centered roughly on the right edge of the ball.

5. Click on the "Track Continuously" button (the button farthest to the right). Watch as the sequence of images is loaded and the Threshold Tracking AOI centers itself on the right edge of the ball.

6. Take a look at the results that Tracker saved to a file. Click on the “Tracking→Result File...” menu item. A dialog box that lets you select the file in which to store the results will pop up. You can see that Tracker has already put the results in the default file named “results.txt”. Click on the “View” button to load that file so you can see what Tracker found. You will see a table that shows the X and Y position of the right edge of the ball (measured in pixels by default) verses time (measured in frames by default).

That’s all there is to basic tracking. Things were slightly simplified here because there was not much needed in the way of image processing and the default tracking parameters were acceptable. Here is a more complicated example that continues from where the previous example leaves off:

1. Close the result file, cancel the “Result File” dialog box, rewind back to the first image, and drag the Threshold Tracking AOI back to center on the right edge of the ball.

2. Imagine that instead of a nice, clearly defined ball (designed to be easy to track), you are working with a poorly defined particle that needs some image processing to find the edge. Click on the “Process→Edge Detect...” menu item. A dialog box with different types of edge detection algorithms will pop up. The default “Prewitt” will work, so click on OK. Notice that the image is processed with this algorithm inside of the AOI immediately.

3. Next apply the skeletonizing image processing algorithm by clicking on the “Process→Morphological...” menu item, selecting the “skeleton” algorithm, and clicking OK. Notice that the algorithm executes immediately on the image inside of the AOI. There is now a nice, sharply defined line in the AOI that shows the position of the edge.

4. Imagine that Tracker needs to apply those two image processing operations for each frame in order to correctly find the edge. Instead of applying the operations immediately as we have been doing, we need to tell Tracker to do them automatically when it tracks. This is done with the “AOI→Process Sequence...” menu item. Notice that the dialog that pops up already has a “Threshold - Simple” operation in the sequence. This is normal for Threshold Tracking. We want to add our two image processing operations to the sequence. This is done by clicking on “Add”, “Edge Detect”, and “Add”, “Morphological”.

5. The operation’s subtype and other parameters (i.e., “skeleton” for the morphological operation) default to the last parameters you selected, so the correct values are already inserted into the sequence list. But, you could click on “Edit” at this point to change the operation subtype or other parameters. As an example, click on the “Threshold - Simple” item in the sequence list to highlight it, then click on “Edit”. A dialog box will pop up showing the parameters of the threshold operation required as the last item in the sequence when threshold tracking. The parameters are set to do a “simple” threshold that sets all of the pixels with an intensity below 128 to black (level 0) and all of the pixels above 128 to white (level 255). Click on OK twice to close the dialog boxes. That sequence of 3 image processing operation will be performed each time the AOI is updated while tracking.

6. Next, create a second AOI to track the left edge of the ball. Clicking on the “AOI→New→Tracking→Threshold Tracking” menu item again does this. A second Threshold Tracking AOI will appear next to the original AOI. Notice that the new AOI is bolder, and the first AOI is slightly dimmer. This is to let you know that any settings you change now will affect the second AOI and not the first. Also notice the two buttons with the arrows and green boxes have become active. These are the “Next AOI” and “Previous AOI” buttons that let you select the AOI you want to work with. Try them out and watch what happens to the AOIs.

7. Make sure the new second AOI is highlighted, and drag it over to the left edge of the ball. The arrow on the AOI still points to the right, but we want it to point left to track the left edge. Dragging the head of the arrow over to the left can change it, but it is more precise to do this from a dialog box. Click on “AOI→AOI Options→Threshold Parameters...” to bring up the Threshold Tracking settings dialog box. Click on “Left”, then “OK” to set the AOI to track the left edge.

8. Click on the “AOI→Process Sequence...” menu item. Notice that the two image processing operations that we added to the first AOI are also there for the second AOI. This is because when you create a

new AOI with one of the same type highlighted, the first AOI is duplicated as much as possible. Click “OK” to close the dialog.

9. Click the “Track Continuously” button to begin tracking, and look at the results just as you did before. Now there are position results from both AOIs verses time.

Users Guide

The users guide is intended to explain Tracker concepts. For installation instructions, see the Getting Started section. For detailed explanations of all Tracker menu items, buttons, and dialog boxes, see the Reference Guide.

Tracking procedures are designed to be mostly the same regardless of the source of the images or the type of tracking that is to be performed. This means that concepts from the extended example above should apply to most tracking situations. However, the wide variety of operations that can be combined in a “process sequence” gives Tracker the ability to do much that is not immediately obvious. It is helpful to understand the basics of how Tracker works before trying to come up with a sequence of processing steps needed to solve your particular tracking problem.

Image sources

Tracker is designed to track objects from a series of images. These images can be from a numbered sequence of disk files, or they can be acquired with a frame grabber from a video source.

Files

The first image in a sequence of disk files is loaded using the “File→Open” command on the menu (or the equivalent button on the button bar, see fig. 2).

For Tracker to load a sequence of disk files, the files have to be stored in a format that Tracker can understand, and the files must be named with a convention that Tracker can understand.

The naming convention is of the form “basenameNNN.TTT”, where NNN is an integer number, and TTT describes the file format. For example: “droplet42.tif”. Only one image per file is allowed, and the numbers in the file names must be sequential for Tracker to treat the files as a sequence. The first part of the file name (“basename”) can be anything.

Tracker always saves images in the tif file format, but it can read images in any of the following formats: tif, gif, tga, dxf, eps, img, jpg, pcx, png, wmf, wpg, PhotoCD and bmp. Several of these formats (especially tif) support optional features or subtypes that might not be supported by Tracker. Tracker will automatically try to recognize files in different formats, and will generate an error message if it fails to understand the file.

Grabbed images

To grab images form a video source (such as a VCR), you must first select the type of video source from the “Transport→Device Selection” menu. The first time you do this, you will be prompted (if necessary) for the serial port parameters necessary for Tracker to send commands to the device. These setting are remembered, so you do not have to enter them again. If you ever need to change the settings, select the device first (as above), then select the “Transport→Device Communication Setup...” menu item.

A single frame from the video source can be grabbed manually with the “Transport→Grab Frame” menu item (or the equivalent button that looks like a still camera; see fig. 2). It is inconvenient to keep clicking this button while manually positioning or adjusting a transport, so Tracker also provides a “Transport→Grab Continuously” menu item (and the equivalent button that looks like a movie camera) that continuously grabs frames and displays them on the screen. Select this menu item or button again to stop grabbing continuously. “Grab Continuously” can also be used for just watching a tape on the computer screen.

Note that most VCRs have a “local/remote” switch that controls whether the VCR will accept commands from the computer interface or the buttons on its front panel. If you are having trouble manually positioning the tape to the correct starting time because the VCR seems to be ignoring the front panel controls, check to make sure the local/remote switch is set to “local”. If Tracker is complaining that the VCR is not responding to its commands, check to make sure the local/remote switch is set to “remote”.

The Matrox Meteor frame grabber that Tracker uses has some optional settings for controlling digitization. These options are under the “Transport→Device Specific Commands” menu item. S-Video or NTSC input can be selected (these are two different input connectors on the frame grabber board). The brightness and contrast of the digitization process can be adjusted from the dialog box (fig. 4) displayed by selecting the “Transport→Device Specific Commands→Digitizer Settings...” menu item.

Viewing images

Once an image is acquired (either from disk or the frame grabber) it is displayed in the main part of the Tracker window. At the bottom right of the Tracker window is a status display that always shows the position of the cursor within the image, and the color and intensity components of the pixel under the cursor.

By default, position is measured in terms of pixels, with the upper left-hand corner of the image being the origin. The first (or X) component of position is zero on the left of the image, and increases to the right. The second (or Y) component of position is zero at the top of the image, and increases in the down direction.

For color images, the status display consists of the red, green and blue components of the color, as well as the intensity component (for which Tracker takes to be the average of the red, green and blue components). For grayscale images, the status display consists of just the intensity. All of these components are eight bit values that range from 0 to 255 (for example, in a color image, black consists of the RGB components 0 0 0, and white would be 255 255 255).

Transport controls

Regardless of whether the images come from disk files or the frame grabber, Tracker simulates a set of VCR-like transport controls (buttons for Play, Stop, Rewind, etc.) that can be used to step through the sequence of images. These controls appear both as buttons on the button bar and on the menu items under the “Transport” menu item. There is also a “Cancel” button that can be used to cancel an outstanding RS-232 command to a device that is stuck waiting for the device to respond.

Zooming

Tracker can change the magnification of the image being viewed by zooming in or out. This can be done with the “View→Zoom In” and “View→Zoom Out” menu items, or the equivalent buttons on the button bar (the yellow magnifying glasses with the “+” and “-” inside of them). The “View→Zoom Original” menu item sets the display back to the original 1:1 zoom level.

Panning

If the Tracker window is not big enough to display the whole image at the current magnification, scroll bars will appear at the edge of the display. Clicking the mouse on these bars allows the image to be scrolled so that different parts of the image can be seen. Note that the presence of scroll bars significantly slows down the speed at which the frame grabber can write data to the screen. If the live display seems sluggish and the scroll bars are present, try a smaller zoom factor to eliminate the scroll bars (and speed up the display).

Scaling

The position of points within the image does not have to be measured in terms of pixels, and time does not have to measure in terms of “frames” (Tracker normally considers each image loaded from disk or each images grabbed to be a frame, typically $1/30^{th}$ of a second). Tracker can convert the position of pixels within an image into the engineering units of your choice. This is done from the “Pixel to Engineering Unit Scaling” dialog box (see fig. 3), which appears when the “View→Image Scale” menu item is selected.

Position

When this dialog is displayed, two markers (the “origin” and the “target”), joined by a line, are superimposed on the image. The origin marker is shaped like a large “V” aligned with the X and Y axes, and by default is in the upper left corner of the image. The target is shaped like a small plus sign with concentric circles, and by default is in the lower right corner of the image. The dialog box shows the position of the origin and target in terms of both pixels and engineering units (which by default, are the same values).

The markers can be dragged to different positions with the mouse. The position of the origin marker is always considered to be the (0, 0) point when reporting position, and the position of the target is reported relative to the origin (and therefore, might include negative numbers). The position of the target in engineering units is updated as the markers are moved, and has the feature that a new value can be entered to replace the calculated value. If this is done, a scale factor is calculated that will map the number of pixels in the image into the number specified for engineering units. Alternatively, a scale factor can be entered directly, and the new values for engineering units will be calculated. In either case, the scaling and origin position will be used for all future position measures that Tracker makes.

The origin can also be rotated to align the X and Y axes with image features. This can be done by typing in an angle in the “Rotation” box, or by dragging the X axis of the origin (click and hold near the small “x” that indicates the X axis).

In addition, the sign of the Y axis can be inverted (for example, to make the measurements correspond to a positive X is right, positive Y is up system). This is done by checking the “Invert Y axis” box.

Time

Tracker will, by default, report time in terms of “frames”, which is really just a count of the number of images Tracker processes. The “Time Scale” box can be used to scale this frame counter into engineering units. A video frame is usually 1/30th of a second, so the Time Scale value could be changed to 30 f/s to have Tracker report time in terms of seconds. Film shot at standard cinematic rates is usually 1/24th of a second per frame. In this case, making the Time Scale value 24 f/s, yields time in seconds.

Another option for time reporting is available by checking the “Use real-time clock” box. This causes Tracker to report time in terms of the elapsed real time according to the system clock on the computer. This is intended to be used when Tracking from a live video source (in this case, the time between frames is not fixed; it depends on the speed of the computer that Tracker is running upon and the amount of image processing Tracker is required to do for each grabbed image). It may be possible to get useful data from Tracker using this technique with a VCR or other video source that is not supported by Tracker. The VCR can be manually put into “play” mode, and Tracker will run unsynchronized to the tape. There will not be a one-to-one correspondence between the images Tracker sees and the frames on the tape, but this may not be needed in all applications.

Areas Of Interest (AOIs)

Tracker works with one or more subsets of the full image. Each of these subsets is called an “Area Of Interest” or AOI. An AOI is used to select an area of the image to which image processing operations will be applied, in which image measurements will be taken, and in which to search for objects to track.

Manipulating AOIs

The “AOI” menu contains all of the commands for manipulating AOIs. Many AOI parameters (such as their size and position) are changed with the mouse.

Creating and deleting AOIs

New AOIs are created with the “AOI→New→...” menu items. There are currently nine types of AOIs, divided into two categories: Tracking AOIs and Tool AOIs. Tracking AOIs are AOIs that save useful information about each frame of a sequence of images. All of the Tracking AOIs (except the Character Recognition AOI) can move themselves around the screen to follow objects that move from frame to frame. The Tool AOIs are useful for interactive manipulation and measurement of the image (for example, you might use a Tool AOI to experiment with different image processing algorithms to determine which ones work best for your image). Tool AOIs are also useful for masking-off unwanted areas of the image so that the Tracking AOIs will not process data in the masked area. Tool AOIs never change position on their own.

A single AOI can be deleted with the "AOI→Delete" menu item. The "AOI→Delete All" menu item will delete all of the AOIs in an image.

Next and Previous AOIs

If there is more than one AOI in an image, most menu items will apply to only the "currently selected" one. The currently selected AOI is drawn with thicker lines and might (depending on the type of AOI) show more details than the other AOIs. The "AOI→Next" and "AOI→Previous" menu items (and the two corresponding tool bar buttons that look like green boxes with arrows) will cycle through all of the AOIs in an image, selecting each one in turn. Information about the currently selected AOI will appear in the status bar at the bottom of the Tracker window.

Moving and sizing AOIs

Dragging an AOI's control points around with the mouse changes the size or shape of the AOI. Dragging the mouse starting in a position away from any of the control points will move the entire AOI. A control point is typically a corner of an AOI.

AOIs can also be moved and sized with the keyboard. This is most useful for fine tuning an AOI's position. The cursor movement ("arrow keys") will move an AOI around if no control point is highlighted. Each control point of an AOI can be highlighted in turn by repeated pressings of the "Tab" key. When a control point is highlighted, the cursor movement keys will move the control point. The escape key ("Esc") is used to remove the highlight from all the control points, so that the cursor keys can be used to move the AOI (clicking the mouse away from any control points also removes the highlight).

Some of the AOIs (the "tracking" AOIs) move themselves around from one frame to the next to track objects. To help them follow quickly moving objects, they take the objects "velocity" into account. This is done by calculating how far and in what direction the object moved compared to the previous frame, and (assuming that it is likely that the object will continue at that speed and direction) moving the AOI by that amount as soon as the frame is advanced (prior to searching for the object to track). This does not work for the first frame of a series, since there is no previous frame to use to calculate the velocity. However, it is possible to set the initial velocity for an AOI with the mouse by holding down the shift key and pressing the left mouse button on one of the AOI's control points, then dragging a line to set the velocity. When the frame is advanced for the first time, the whole AOI will move by the length of line and direction indicated.

AOI Options

The Tracking AOIs all have a number of optional settings that control how the tracking is done. Selecting the "AOI→Options→..." menu items will bring up dialog boxes that allows these settings to be edited. Each different type of AOI may have different options.

If a new AOI is created that is the same type as the currently selected AOI, it will be created with a copy of all the current AOIs options. The new AOI's position will be offset slightly so that it is obvious that there is a new AOI, but it will be otherwise identical to the original AOI.

Image Processing

Tracker includes a large number of image processing operations that can be applied to enhance or extract information in the images. Image processing operations are divided into about 10 major categories that are listed under the "Process" menu item. Selecting something from the "Process" menu (for example "Process→Threshold...") will bring up a dialog box that lets you select the sub-type of the image processing operation (for example, a simple threshold verses a band-pass threshold), as well as all of the parameters needed to apply the operation (i.e., threshold levels and transition points). Clicking the “OK” button in this dialog will immediately apply the selected image processing operation to the currently select AOI.

When tracking, a sequence of image processing operations may be required to be applied to each AOI after every frame is advanced. The sequence is specified for a particular AOI by selecting the “AOI→Process Sequence...” menu item. This brings up a dialog box (see fig. 15) that allows a sequence of image processing operations to be defined and manipulated. The “Add” button adds a new image processing operation to the list (using the default parameters for that operation). The “Edit” button allows the parameters of the operation to be changed. The “Delete” button deletes the currently highlighted operation. The “Move Up” and “Move Down” buttons are used to change the order of the operations in the list. The “Load” and “Save” buttons are used to store and recall the entire list to a disk file (so that you don’t have to regenerate a complicated list every time you use Tracker).

Updating AOIs

The image processing sequence can be applied to one AOI by selecting the “AOI→Update” menu item. After the image processing operations are applied, the AOI will do whatever it normally does to update itself (for instance, a Histogram AOI will recalculate and display the histogram; a Threshold Tracking AOI will search for the farthest pixel in an object in a particular direction, and move the AOI to center on that pixel). In fact, “Updating” an AOI causes the AOI to do everything it will do when tracking with 3 exceptions: the frame is not advanced first, only the current AOI is processed, and no results are saved to the result file. This is useful for making sure the AOI will behave as expected once tracking starts.

The Tracking AOIs

The Tracking AOIs are designed to process and extract information from a sequence of images. All except the Character Recognition AOI will move themselves around attempting to follow object identified in the images.

Manual

The Manual Tracking AOI requires a person to tell the AOI where it should center itself when it updates. This means that for each Manual AOI in each frame of the image sequence, somebody manually must click the mouse to specify the object being tracked (and this position is recorded in the result file). This is useful if there are only a few frames to track, or when all else fails.

Threshold

Threshold Tracking AOIs convert the image into a binary form (just black and white), and searches for the farthest white pixel in a specified direction. The location of this pixel is recorded in the result file, and the AOI is moved so that it is centered on the location (this causes it to follow the edge of an object from frame to frame). This is the most common type of tracking.

Snake

The Snake Tracking AOI allows the user to define a cubic spline curve that will interact with the image and move itself around (called “slithering” in snake parlance) to search for image features. This is useful for tracking an interface or edge, and is also sometimes helpful when objects overlap. In addition, the Snake Tracking AOI can optionally record the position of all of the points along its curve, so it is also used when the shape of an object (instead of just its position) must be recorded.

Pattern Match

The Pattern Match Tracking AOI uses correlation based techniques to find the positions of best matches of a given example image to each of the frames in a sequence. For example, it could be used to identify all of the parts of the image that look like a given picture of a droplet. Positions where the correlation exceeds some specified threshold are recorded in the result file.

Character Recognition

The Character Recognition AOI finds the best match between the characters in an image and a pre-defined alphabet (usually just the digits 0 through 9). The recognized characters are stored to the result file. This is used to extract time stamp or instrumentation information from the image (for example, some images have a digital clock near the objects of interest to show when the data was generated). The Character Recognition AOI does not move itself around to track objects (the characters are assumed to always be in the same location of the image).

Tool AOIs

Tool AOIs are designed to most often be used interactively (rather than being used for tracking). They are useful for experimenting with image processing operations to find the ones that work best with a particular set of images. However, they are sometimes useful with thresholding operations while tracking in order to mask off a part of the frame that might confuse one of the tracking AOIs. None of the tool AOIs move themselves around from frame to frame to track objects.

Rectangular

Rectangular AOIs are simple rectangles. They are the most common AOIs used to experiment with images processing operations.

Polygon

Polygon AOIs can have any polygonal shape, and are therefor very useful for masking off oddly shaped parts of an image.

Histogram

The Histogram AOI is a rectangular AOI associated with a separate window that shows a histogram of the distribution of pixel intensities or colors within the rectangle.

Line Profile

The Line Profile AOI is a line associated with a separate window that shows the sequence of intensities of colors along the line. This is very useful when you need to determine what a good threshold level would be for a particular image.

Image Processing

Operations

Tracker supports a large number of image processing operations. These operations are used to make the image details that are important stand out so that they may be tracked. Operations may be performed on an AOI interactively (by selecting them from the "Process" menu), or they may be included in an image processing sequence to be applied to each frame while tracking.

Un-do and Re-do

Finding the right combination of image processing operations for a set of images often involves a lot of experimentation. To make this easier, Tracker supports "Un-doing" and "Re-doing" of interactive image processing operations. After an image processing operation is performed, the "Process→Undo" and

“Process→Redo” menu items will be enabled (along with the corresponding tool bar buttons that perform the same functions).

The “Un-do” function reverts the AOI to the state it was in before the image processing operation was applied. If a series of image processing operation was performed, repeatedly pressing “Un-do” will step back through the entire series. Pressing “Re-do” will step forward in the series, re-applying the operations as they were performed before. If you are at the end of a series (i.e, you did an image processing operation, but did not un-do it), the “Re-do” function can be used to quickly repeat the operation. This is useful for repeatedly applying image processing functions such as “smoothing”.

Tracking

Tracker performs the following steps when tracking:

- The transport is advanced to the next frame
- Each AOI is processed, in order, by:

    Applying the sequence of image processing operations

    Finding the tracking results for this AOI (edge of the object, etc)

    Moving the AOI so it centers on the object found

- The accumulated tracking results for this frame are stored to a disk file

This tracking process can be performed for one frame by selecting the “Tracking→Track 1 Frame” menu item (or by clicking on the equivalent button on the button bar: the button with the “T” and the straight arrow, see fig. 2). This is primarily used for making sure that everything is set up properly before tracking a series of images.

Starting and stopping

The main purpose of Tracker is performing the tracking steps above for an entire series of images. This is initiated by selecting the “Tracking→Track Continuously” menu item (or by clicking on the equivalent button on the button bar: the button with the “T” and the arrow that loops back to point at the “T”, see fig. 2). By default, the tracking process will repeat until the transport can no longer advance to the “next frame”, or until it is stopped manually (by selecting the “Track Continuously” button again). Tracking can be stopped and restarted as many times as is necessary.

There are several options available to determine how Tracker determines the “next frame”. They are available by selecting “Transport→Next Frame Setup” from the menu. Tracker can advance the transport by more than one physical frame for each image that is tracked (useful for reducing the number of frames processed for slowly changing phenomena). Tracker can also average several physical frames together to get the image used for tracking (useful for reducing noise with slowly changing phenomena). Finally, Tracker can be set to stop tracking after a fixed number of frames.

Result files

Tracker stores all of its results in text format in a “result file”. By default, this file is named “Result.txt” and is put in the same directory as the Tracker program. You can select a different file for the results, view the result file, and empty the result file (erase anything already in it) from the dialog that appears by picking the “Tracking→Result File...” menu item (see fig. 28). New data is always appended to whatever is already in the result file.

The “Result File” dialog also allows some options to be changed concerning how the data is stored in the file. By default, every time tracking is started a column header is written to the result file, and the first frame tracked is reset to be “Frame number 1”. Both of these behaviors can be suppressed by checking the appropriate option boxes. This is useful when tracking must stop and start in the middle of a series of images. The defaults assume that a new series is begun each time tracking is started.

Reference Guide

The Reference Guide provides a detailed description of the menu interface, buttons, and dialog boxes.

Tracker Operation

For the following description of Tracker operation, see figure 2, below. Tracker displays images inside what is known as its Client Area (viewing area), which is the white area just below the button bar. If the image is too large to be displayed fully inside Tracker viewing area, scroll bars appear, and can be used to scroll the desired part into view. Tracker can be operated fully using the menus, but some operations are done easier and quicker using the buttons on the button bar. All of the buttons have a corresponding menu item; in fact, the buttons may be thought of as shortcuts to the menu items. The status bar is used to display various types of information that Tracker continually puts out. It is divided into three main panels. Panel 1, on the left side of the status bar, displays information about Tracker's current activity, primarily information about the currently selected AOI. Panel 2 displays the cursor position and color and intensity values at the cursor position. Panel 3 is reserved for a count of frames tracked. This is the number that is written into the data file. It is useful during automatic tracking for seeing how many frames have been tracked. The title bar is used for indicating the image file name, when the images are obtained from disk.

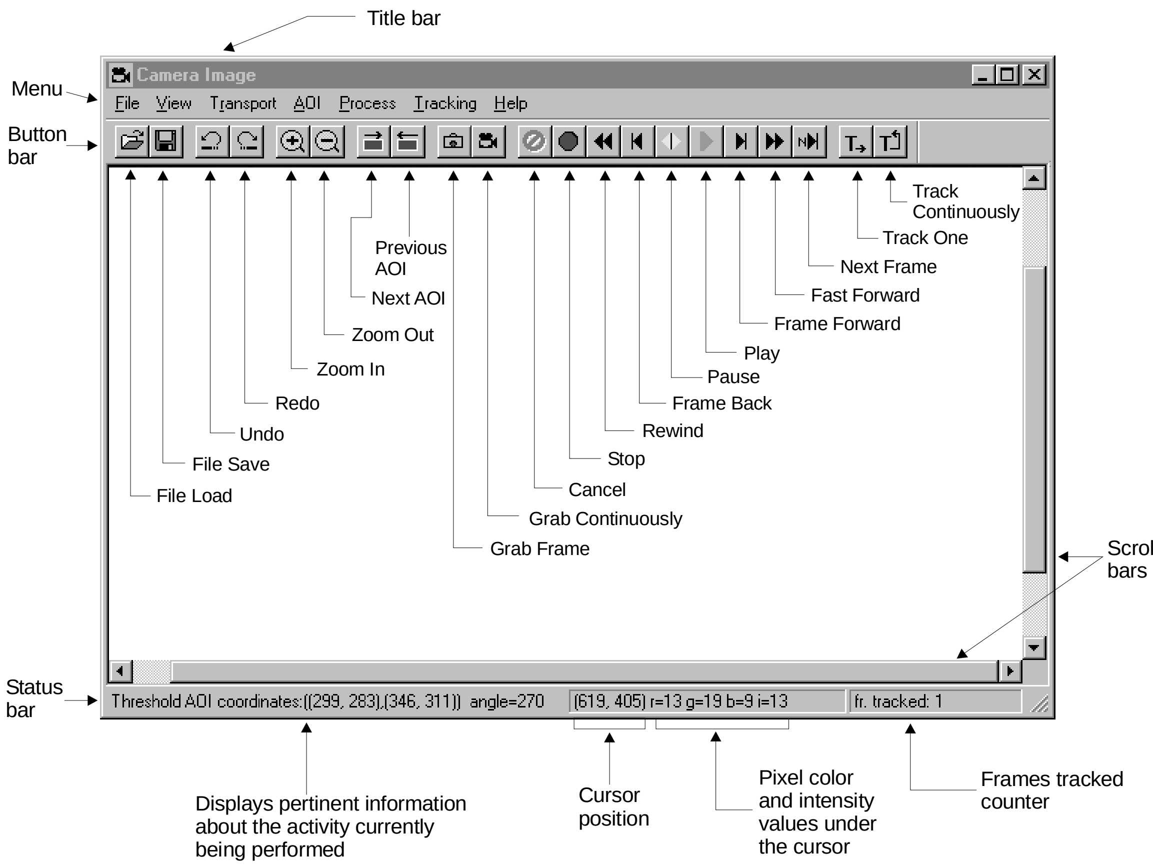

Figure 2. – Tracker options.

File

The menu items under the File menu pertain to loading and saving images to or from disk.

Open – loads image from the hard disk. The image file formats supported are tif, tga, bmp, gif, dxf, eps, img, jpg, pcx, png, wmf, wpg, and PhotoCD. If another device is selected at the time (such as JVC tapedeck) and an image file is loaded from the hard disk, the device selection is changed to Disk Files. This menu item has a corresponding button on the button bar, labeled “File Load” (see figure 2).

Save As – saves the entire image to a hard disk or to any connected drive. The only image file format currently supported is TIFF. This menu item has a corresponding button on the button bar, labeled “File Save” (see figure 2).

Save AOI – saves the portion of the image enclosed by an Area of Interest to disk. The only image file format currently supported is TIFF.

Exit – exits from Tracker.

View

The menu items under the View menu pertain to viewing the image, viewing information about the image, or setting up how the image is to be viewed.

Tool Bar - a toggle which enables or disables the Tracker tool bar (also called the button bar). The Tracker button bar is shown in figure 2.

Status Bar - a toggle which enables or disables the Tracker status bar. The Tracker status bar, located at the bottom of the Tracker window, is used for displaying various messages and tracking information (see figure 2).

Zoom In – zooms the image in 2x each time the menu item is selected. The pixel information in the image is not affected, only the view is. Note the corresponding button on the button bar (see figure 2).

Zoom Out – zooms the image out 2x each time the menu item is selected. The pixel information in the image is not affected, only the view is. Note the corresponding button on the button bar (see figure 2).

Zoom Original – the zoom factor is returned to 1.

Image Scale – this menu item brings up the Pixel to Engineering Unit Scaling dialog box (see fig. 3) through which several items can be set, including spatial and temporal scale factors, coordinate rotation, and coordinate origin. The spatial factor allows one to track in terms of real spatial units such as millimeters rather than in pixels. The temporal scale factor allows one to track in terms of real temporal units such as seconds rather than frames. Coordinate rotation allows rotating the coordinate system so that an off-axis linear motion can be aligned to one of the axes, essentially reducing the 2-D motion to 1-D motion. And finally, the coordinate origin can be set to any location in the image, including the object position in the first frame. This may be useful for tracking an object relative to some stationary point in the image.

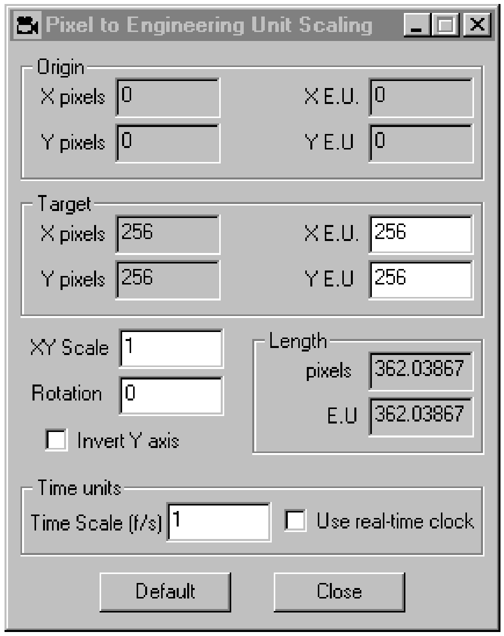

Figure 3 – Pixel to Engineering Unit Scaling dialog box.

Origin – displays the coordinates of the origin of the scale tool. The left side shows the coordinates in terms of pixels and right side in terms of the engineering units determined by the XY scale. The default origin is in the top left corner of the image and the positive Y direction is "down".

Target - displays the coordinates of the target end of the scale tool. The left side shows the coordinates in terms of pixels and right side in terms of the engineering units determined by the XY scale. The default target position is in the bottom right corner of the image.

XY Scale – the spatial scale factor. It is set in two ways. One is by simply typing it in. The second is by typing in the Target engineering units and letting Tracker calculate the scale factor based on the dimension of the scale tool (the difference between the origin and target positions).

Rotation – the angle of rotation of the scale coordinates. This factor can either be entered by hand or it is interactively updated as the user rotates the scale tool with a mouse. (This is done by clicking and dragging the small "X" displayed on the X axis.)

Invert Y Axis – checking this option inverts y-axis causing y values to increase in the upward direction.

Length – reports the length of the scale tool in terms of pixels and engineering units.

Time Scale (f/s) – the frame rate in terms of frames per second at which the image sequence was recorded.

Use real-time clock – changes what is written to the output file in the first column from frames to system clock readings. This option is intended for use with real-time or near-real-time tracking when live camera or unsupported tapedeck is used. The user may select this option and let the tracking proceed at whatever speed the computer can handle. Real-time tracking (30 f/s) might be approached with very fast computers.

Image Info – displays information about the image being displayed, such as the size and number of color planes.

Transport

The menu items under the Transport menu set up and/or control the image input devices (transports).

Device Selection – selects the image input device from a list of devices that Tracker can communicate with. Selection of one of the devices sets up a default communication configuration for that device including the type of image acquisition and serial communication parameters. Both of those configurations can be changed with the next two menu items.

Device Specific Commands – this menu item contains additional commands that are specific to the selected device. For the video devices, this includes options for selecting between NTSC and S-VIDEO input and an option to change the digitizer settings.

***S-Video (YC)* –** (for video devices only) – selects S-Video as the input to the frame grabber (digitizer).

***NTSC* –** (for video devices only) – selects NTSC as the input to the digitizer.

***Digitizer Settings* –** (for video devices only) – brings up a dialog box (see figure 4) through which the frame grabber brightness and contrast can be changed.

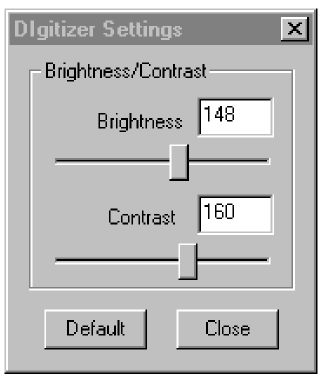

Figure 4. – Digitizer Settings dialog box.

Brightness – changes the image brightness.

Contrast – changes the image contrast.

Device Communication Setup – brings up a dialog box (see figure 5) through which the serial communication parameters are selected or changed. Each device (that is controlled with a serial interface) has serial port parameters set through this dialog box. This information is saved and becomes part of the device profile until changed.

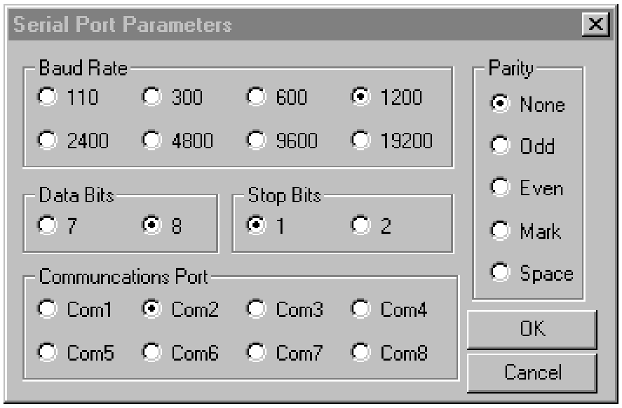

Figure 5. – Serial Port Parameters dialog box.

Baud Rate **–** communication rate in bits per second that the device is using.

Data Bits **–** communication protocol that the device is using.

Stop Bits **–** communication protocol that the device is using.

Parity – communication protocol that the device is using.

Communication Port – serial port to which the device is connected.

The following group of menu items control the operation of the selected device and correspond to a group of buttons on the button bar. The selection of one of them or clicking the corresponding button on the button bar communicates the command to the selected device and causes the device to take the appropriate action.

Stop – instructs the device to stop any activity.

Rewind – instructs the device to rewind at full speed.

Frame Back – instructs the device to step back one frame.

Pause – instructs the device to pause. Pause should be used instead of “stop” in most instances as stop might cut off image input as well as rendering the device inactive, whereas “pause” maintains image display while temporarily suspending motion.

Frame Forward – instructs the device to step forward one frame.

Play – instructs the device to advance at its standard speed.

Fast Forward – instructs the device to move forward at full speed.

Next Frame – instructs the device to advance by a pre-set number of frames. The number of frames can be entered in the Next Frame Setup dialog box (see figure 6).

Next Frame Setup – brings up a dialog box (see figure 6) through which one can set frame step number, total number of frames to track, and number of frames to average.

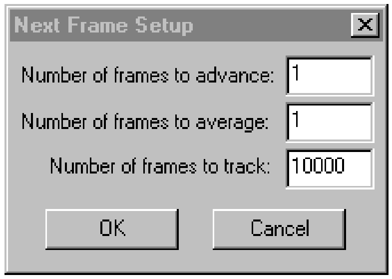

Figure 6. – Next Frame Setup dialog box.

Number of frames to advance – sets the step size in terms of frames. This item works in conjunction with the Time Scale set in the Scale dialog box (fig. 3). For example, if the Time Scale is set to 30 f/s and number of frames to advance is set to 2, then Tracker will take steps of 0.066 sec. (or 2 frames worth).

Number of frames to average – sets a number of frames to average. This is a moving average where one frame is subtracted from the average as the new one is added in.

Number of frames to track **–** causes tracking to stop after this number of frames.

Grab Frame – captures a single video or film frame. Note the corresponding button on the button bar which looks like a still camera, see figure 2.

Grab Continuously – sets the frame grabber into a live capture/display mode. Note the corresponding button on the button bar which looks like a movie camera, see figure 2.

AOI

The menu items under the AOI (Area of Interest) menu deal with setting up, managing, and deleting AOIs. An AOI is a subset (region) of the image in which all operations take place. The AOIs are split into two categories: ones used for tracking (under Tracking menu) and ones usually used for analysis (under the Tools menu). All tracking AOIs, except for Character Recognition, move with the object being tracked whereas the analysis AOIs are immobile. The analysis AOIs are typically (but not always) used with a single frame analysis and the results of any analysis is typically (but not always) only visual; no information is written to a file.

New – creates a new AOI of one the following types.

Tracking – the selections from this menu create an AOI that is used for tracking. Selection of a particular type of AOI actually defines a "tracking method", which then uses the selected AOI. The tracking method is what is described below. A default version of the AOI is drawn on the screen and is typically modified and moved using the mouse. Some image enhancement is typically performed to help with the tracking. The moving AOIs (all tracking AOIs except the Character Recognition) also have a velocity function associated with them with predicts the forward motion and makes a better judgement on the placement of the AOI in the next frame (to increase its odds of finding the object inside the AOI).

Manual Tracking – a rectangular AOI used for manually locating the position of the object using the mouse. Once the tracking process is started, the user clicks left mouse button on the object to identify its position and Tracker performs all the advancement of frames, image processing, and storing of the object position coordinates to a file.

Threshold Tracking – this tracking method tracks the edges of objects by performing an image threshold operation and then searching for the first occurrence of non-black pixels in a given direction. The first point found is called the track point. Some image processing is typically performed to reduce noise before the thresholding operation. The AOI moves with the object by automatically centering it self on the track point in each frame. The Threshold Tracking Parameters dialog box is shown below.

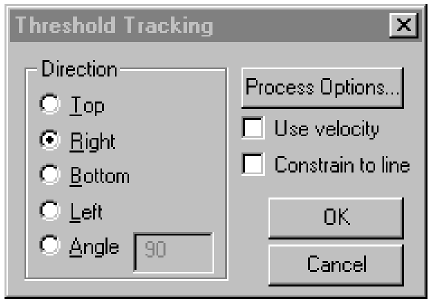

Figure 7. – Threshold Tracking dialog box.

Direction – selects the tracking direction; i.e., this angle determines the direction in which AOI is searched. For example, if the direction selected is "Right", or angle of 90 degrees, the AOI is searched from right to left for the first occurrence of a non-black pixel.

Process Options – brings up the threshold processing options dialog box, shown in fig. 17, from which a threshold method can be selected.

Use velocity – enables prediction of the AOI's position in the next frame based on the current velocity of the object.

Constrain to line – constrains the AOI repositioning during tracking to a line. This may be useful in forcing an AOI to follow a predetermined path without being sidetracked. When this option is selected (and after OK in the dialog is clicked), a dashed line appears across the image. This line can be rotated by dragging the arrow tip of the AOI angle selection arrow.

Snake Tracking – the snake tracking method tracks an object using an "active contour line". This is a flexible cubic spline curve that interacts with features in an image to move itself around ("slithering", in snake parlance).

The snake algorithm works by moving the control points of the spline curve around, evaluating the "energy" at each position of the control points, and moving the curve to the spot where "energy" is lowest. The energy function determines how the snake will interact with features in the image and move. The energy function is a weighted sum of contributions from the various snake parameters (such as, how close are the image pixels under the snake curve to the desired ideal intensity). The weights determine the relative importance of the snake parameters. There are complex interactions between many of these parameters, so it is best to experiment with different values to try to find something that works well for your images.

The tracking results stored while snake tracking can be the average of all points along the snake line (for centroid-like measurements) or it can be the position and colors of the actual pixels under the line. Snakes may be closed into a loop or open-ended, and control points may also be pinned down to keep them from slithering.

The snake is created, moved, and resized using the mouse. To create a snake, control points are added by clicking the right mouse button on the image. Double click the right mouse button to make the snake a closed loop. If the snake is closed, clicking the right mouse button deletes the current snake and starts a new snake. A control point is moved by dragging it with the left mouse button. The control points normally move themselves around as the snake seeks the minimum energy position, but at times it may be desirable to pin down a control point to prevent it from moving. This can be done by holding down the Control key and clicking with the left mouse button on a control point. When the control point is pinned down, its designation changes from a cross to an X. Repeating this process returns the control point back to a movable control point. The Snake Algorithm Parameters dialog box is shown below.

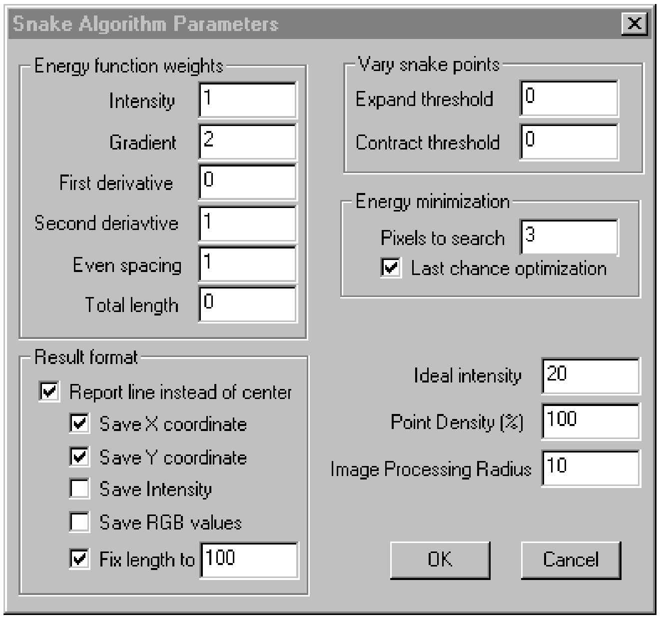

Figure 8. – Snake Algorithm Parameters dialog box.

Energy function weights **–** these parameters control how the snake moves. The energy function is the weighted sum of the terms and as such, their values are only important relative to one another. They are normalized such that their default contributions are approximately one.

Intensity – sets the weight given to the "Ideal intensity" (specified elsewhere in the dialog). It evaluates the difference between the "ideal intensity" and the intensity of the image pixels under the snake, and tends to make the snake seek out parts of the image close to the ideal intensity, and avoid parts of the image that are farther from the ideal intensity.

Gradient – evaluates a gradient along a short line (7 pixels) perpendicular to each point of the snake. The energy decreases when gradient is large. A larger weight for this term causes the snake to seek out edges. Setting this weight to zero avoids the lengthy calculations involved in computing the gradients and greatly speeds up snake calculations (useful if you don't care about edges).

First derivative – increases the energy based on how bent the snake is, but it is dependent on the bending direction and therefore is not of much use. It’s probably best to keep it at 0. This term depends only on the shape of the snake and has nothing to do with the pixels in the image.

Second derivative – increases the energy when the snake has sharp bends, but it is not dependent on direction, so it is much more useful than the first derivative. A large weight here tends to make a closed snake want to be circular and tends to make a non-closed snake want to be a straight line. This term depends only on the shape of the snake and has nothing to do with the pixels in the image. Unfortunately, there are other interactions that affect this term (like the total size of the snake - which has effects on many of the terms), so it is hard to predict exactly what effect this term will have. It may be useful to help keep the snake from forming small loops that have little to do with image.

Even spacing – controls how important the snake thinks it is to have a constant distance between the snake's control points. A larger value here keeps the points from getting all bunched together. This term depends only on the shape of the snake and has nothing to do with the pixels in the image.

Total length – this weight is designed to offset the tendency of many of the other terms to shrink the snake down to nothing. It may be useful to put a small value here if the snake is shrinking too much (or possibly to get the snake to expand to fill a cavity). This term depends only on the shape of the snake and has nothing to do with the pixels in the image.

Vary snake points - these options will automatically introduce (or delete) control points along the snake if the existing points get too far apart (or too close together). This term depends only on the shape of the snake and has nothing to do with the pixels in the mage.

Energy minimization - these parameters control the optimization process used to minimize the energy to find where the snake should move. Basically, the optimization process moves each control point, one at a time, evaluates the energy at the new points, and leaves that control point there if the energy is lowered. The "pixels to search" box selects the number of pixels to initially move the control point. It defaults to 3, so the energy is evaluated with the control point moved to the eight combinations possible by moving it 3 pixels (up, down, right, left, and the four diagonal positions). If any of those positions are better than the current position, the control point is moved to the better spot and the process continues with the next control point. If all the positions are worse, the process is repeated moving only 2 pixels. If that fails, it tries again moving only one pixel. If that fails, it leaves the control point where it was and goes on to the next control point. If all the control points fail to move, the snake decides this is a good spot and stops (unless last chance optimization is checked – in this case all the control points are checked again with the "pixels to search" value doubled [to 6 in the default setup]). This algorithm works fairly fast and usually finds snake locations that seem close to the optimum.

Ideal intensity – sets the pixel intensity level that the snake searches for and works in conjunction with the Intensity weight.

Point Density (%) – sets the number of points along the snake that get evaluated. Setting it to a low number speeds up snake calculations, setting it to a high number increases accuracy. Even when it is set to “100%”, not all the points along the snake are guaranteed to be evaluated (there are complex interactions with snake size and the distance between control points that affect this).

Image Processing Radius **–** sets the size of the processing AOI around the snake. This number is how much bigger the processing AOI will be than the snake’s bounding rectangle. The processing AOI is a rectangular area outside of the snake in which the image processing algorithms are performed (before the snake algorithm is applied). If your snake is moving a lot, the default value of 10 pixels is probably too small. The value doesn't matter if you are not doing any image processing.

Pattern Match Tracking – the pattern match tracking is a correlation-based method in which the algorithm searches for a best match of a template inside a search region. The template can be updated continuously or the same one can be used repeatedly. The search region moves with the object by centering itself on the point found. A number of options are offered including match acceptance level, multiple occurrence acceptance, and tradeoffs between speed and accuracy. The Pattern Match Tracking Parameters dialog box is shown in figure 9, below.

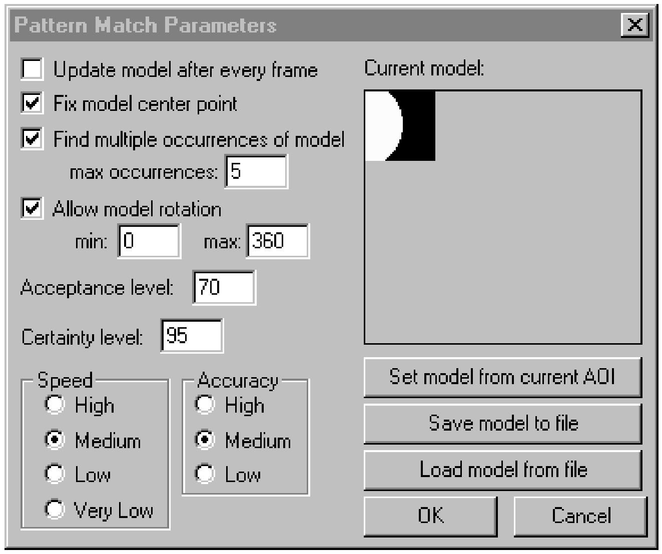

Figure 9. – Pattern Match Parameters dialog box.

Update model after every frame – determines whether the template (model) will be updated after every frame or whether the original template will always be used. Updating the template after every frame allows for changes in the object being tracked. However, sometimes it may be desirable for the template to be fixed to prevent slow divergence from the original.

Fix model center point – when selected, this option fixes (freezes) the report point (the point whose coordinate is reported to the data file) so that it can't be moved accidentally. When de-selected, the report point is allowed to be dragged to a new location. Moving the report point may be useful when the template contains an entire object but the user desires to obtain coordinates at some particular part of the object, even if the object is rotated.

Find multiple occurrences of model – enables the algorithm to report multiple occurrences of the object as long as they meet the acceptance level. The number of occurrences is entered in the box labeled "max occurrences". This may be useful when tracking a number of similar objects within the same search region.

Allow model rotation – enables the algorithm to search for rotated versions of the object. Its possible to limit the rotation range by using the "min" and "max" boxes.

Acceptance level – the acceptance level is the score below which the result of the correlation is not considered a match. A typical match is 80% or higher, depending on the image. For noisier images the level might be set to 70% or less.

Certainty level – the certainty level is the match score above which the algorithm can assume that the object has been found and that it can stop searching.

Speed – the speed parameter allows a trade-off between speed and robustness of search operation (the likelihood of finding the object). The "high" option should be used for good quality images or when the template (model) is simple. The "medium" setting is recommended for medium quality images or for more complex templates. The "low" and "very low" settings should be used only if problems have been encountered at the "medium" setting.

Accuracy – allows sub-pixel accuracy in position estimation. This can be achieved by interpolation of the correlation surface. This is only a theoretical limit and in practice the actual accuracy is affected by the noise in the image and by a pattern in the template (model). Theoretically, the "high" accuracy is ± 0.125 pixel, the "medium" accuracy is ± 0.25 pixel, and the "low" accuracy is ± 0.5 pixel. The accuracy that can be expected realistically for a fairly clean image obtained from a tapedeck at the "medium" level (default), is probably around ± 0.5 pixel, or accurate to within a pixel.

Current model – the current model (template) window displays the image of the template currently in use. The current model can be obtained using one of the three buttons below.

Set model from current AOI – extracts the pixels from the currently positioned AOI and sets this area of the image as the template (model). The template is also displayed in the Current model window.

Save model to file – allows the current template to be saved to a file. A proprietary ".mod" format is used.

Load model from file – loads a previously save template. This template is any ".mod" file.

Character Recognition Tracking – this tracking method is designed for recognition of sequences of numbers and storing them to a file. Currently, recognition of characters other than numbers is not supported. The character recognition is implemented with two different and separate methods: the Neural Net Method and the MIL OCR method. There are strengths and weaknesses associated with each method which will be discussed below. The Character Recognition Tracking Parameters dialog box is shown below.

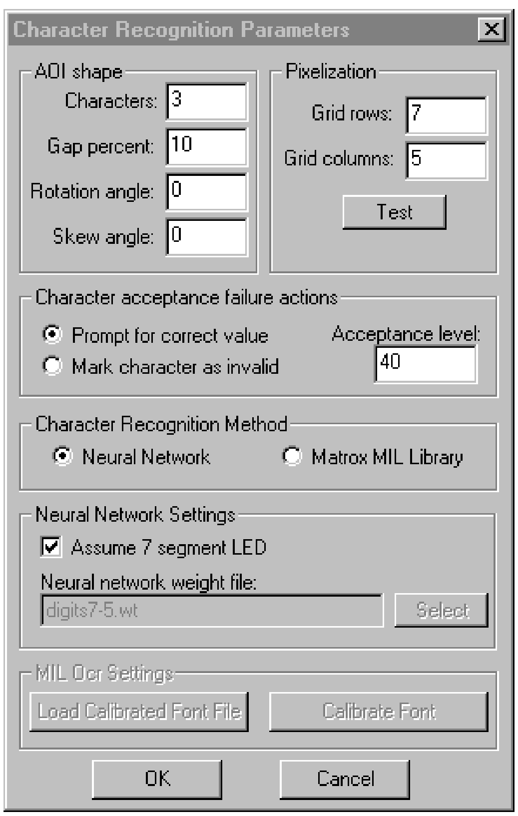

Figure 10. – Character Recognition Parameters dialog box.

AOI shape – these parameters control the shape and orientation of the AOI box. The character recognition box is a special customized box that is divided into a certain number of sub-AOIs. The AOI box encloses the character string with one sub-AOI enclosing each character. The AOI has a special shape capability, which allow it to rotate, skew, and vary the gap between sub-AOIs. A geometric transformation is used to negate the character rotation, skewing, and spacing. See appendix C for more information on defining the shape of character recognition AOIs.

> **Characters** – changes the number of sub-AOIs to match the number of characters in the string. The number of sub-AOIs can be changed in two ways: by typing in the number in this dialog box or by clicking right mouse button on the upper-right (to increase) or upper-left (to decrease) corners of the AOI box.
>
> **Gap percent** – controls the gap spacing in terms of percent of the sub-AOI width. The gap spacing can be changed in two ways: by typing in the number in this dialog box or by dragging the right mouse button on the arrow in the lower-left corner of the AOI.
>
> **Rotation Angle** – controls the rotation of the AOI box. This is useful when the character string runs at some angle relative to the camera. The rotation angle can be changed in two ways: by typing in the number (in degrees) in this dialog box or by dragging the left mouse button on the upper-right corner of the AOI.
>
> **Skew Angle** – controls the skewing of the AOI box. This is useful when the characters in a string are skewed at some angle, which is common on many LED number displays. The skew angle can be changed in two ways: by typing in the number (in degrees) in this dialog box or by dragging the left mouse button on the lower-left corner of the AOI.

Pixelization **–** the pixelization parameter refers to the final resolution to which the sub-AOI will be converted to before the algorithm tries to recognize the character (number). The neural network algorithm is optimized for final character resolution of 5x7 pixels (7 rows, 5 columns) whereas the MIL OCR method uses whatever the pixel resolution of the box is.

Character acceptance failure actions – this parameter selects one of the two available methods for reporting a character that failed the acceptance level test. The algorithm evaluates each character separately and generates a confidence value (1-100). If the algorithm can’t recognize the character, it returns a low confidence value and if this value is below the acceptance level, one of the below actions is taken.

> **Prompt for correct value –** brings up a dialog box prompting the user to input the correct value. This method will be more accurate because it gives the user a chance to insert the correct character even if the algorithm failed, however, user input is required, thus it is not fully automatic.
>
> **Mark character as invalid** – inserts a dash (-) for the invalid character and continues on. This method allows the algorithm to operate completely on its own, but at the expense of some accuracy.

Character Recognition Method – there are two separate character recognition methods supported by Tracker: the Neural Network method and the MIL OCR method. Selection of one or the other enables or disables the appropriate character recognition “Settings” section below.

> **Neural Network –** this method uses a back-propagation neural net technique to recognize numbers. First, the resolution of the numerical digits is reduced down to 5x7 pixels (or optionally, to seven segments that would correspond to a typical LED display). A separate training program (digit.exe, see the appendix) has trained on several sets of reduced resolution numbers, each with a slightly different look, and

generated a set of weights, based on the numbers trained on. This set of weights is used by Tracker to perform the recognition.

MIL OCR – this method uses the Matrox MIL OCR library functions to perform the character recognition. The MIL OCR method uses a calibrated font file (*.mfo) to store all the information about the characters its trying to recognize. In this way it's a bit similar to the weights file that the neural network method uses. To produce a calibrated font file, a tiff file representation of all characters in the font must be made by an external paint package (such as Corel Draw or Photoshop). This tiff file is then one of the inputs to the calibration procedure along with a test string, which produces the calibrated font file. All that the calibration process does is to figure out the size of the character and spacing between characters.

Neural Network Settings – these settings become enabled when the Neural Network option is selected in the Character Recognition Method section.

Assume 7 segment LED – reduces the neural network method to looking only at the presence or absence of the seven segments that comprise the digit. The "seven.wt" file is used for this setting.

Neural network weight file – the neural network method requires a weight file that describes the network. A different file must be used for different values of "pixelization". A file named "digits5-7.wt is provided if the default resolution of 5x7 is used.

MIL OCR Settings – these settings become enabled when the MIL OCR option is selected in the Character Recognition Method section, above.

Load Calibrated Font File – prompts the user for selection of the calibrated font file (*.mfo). This file would have been generated and saved previously using the procedure under Calibrate Font, below.

Calibrate Font – brings up a MIL OCR Calibration dialog box through which the font can be calibrated or the dimensions of the font typed in by hand. The Mil OCR Calibration dialog box is described below.

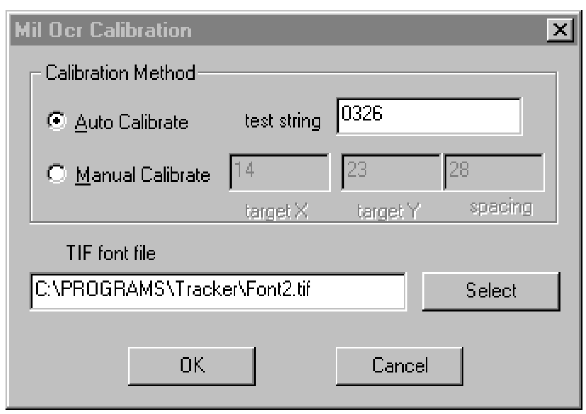

Figure 11. – MIL Ocr Calibration dialog box.

Calibration Method – the calibration of the font involves letting the algorithm know the size and spacing of the target font (text to be read). That is done in two ways: one is to let the computer figure it out (Auto Calibration method) and the second is to just type it in (Manual Calibration method).

Auto Calibrate – lets the computer figure out the size and spacing between the characters based on the test string that is typed in and the font image selected below.

Manual Calibrate – enables the targetX, targetY, and spacing input boxes so the user can input the values in directly.

TIF font image – name and path of the font image file in which the image of the font representation is stored.

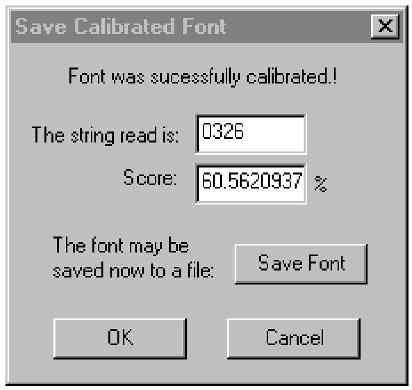

Figure 12. – Save Calibration Font dialog box.

Tools **–** the AOIs found under the Tools menu are designed to be used as single image frame analysis tools, although they can be used across multiple image frames as well. The knowledge gained from the analysis performed with these tools can be used for setting up some of the parameters for performing automated tracking.

Rectangular – creates a rectangular AOI. The AOI can be used to perform image processing trial and error testing. If the results of the image processing tests look satisfactory, then the same image processing steps may be used for tracking.

Polygon – creates a polygon AOI. The polygon is constructed using the mouse in the same fashion as the “snake” by clicking right mouse button to add a control point and double-clicking right button to close the polygon. Image processing can now be performed inside the polygon’s region. A polygon AOI is useful for masking out unwanted objects from the image.

Histogram – brings up the interactive histogram tool for visualizing an area histogram of pixel intensities. The area is described by a rectangular AOI. As the rectangular AOI is moved or resized, the histogram displayed in a dialog box is updated as fast as the computer can keep up.

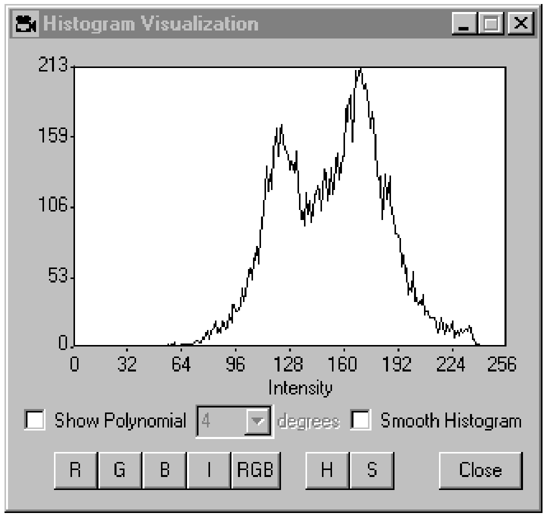

Figure 13. – Histogram Visualization dialog box.

R, G, B, I, RGB, H, S buttons **–** displays an intensity histogram of the color plane labeled on the button. The first four buttons stand for red, green, blue, and intensity, respectively. The next button labeled RGB, displays red, green, and blue histograms all at the same time. The last two buttons display hue and saturation histograms.

Show Polynomial – overlays a polynomial approximation line over the histogram line. The polynomial degree is shown in the drop box to the right. The polynomial may be used as sort of a filtering tool for the actual histogram.

Smooth Histogram **–** performs a 3-point smoothing of the histogram line.

Line Profile – brings up the interactive line profile tool for visualizing pixel intensities under a line. The line can be moved and resized using the mouse. The line profile, displayed in a dialog box, is updated in real time as it is moved and resized. The line profile is useful for graphically illustrating pixel intensities of objects in an image, which can be useful in estimating where threshold levels should be set.

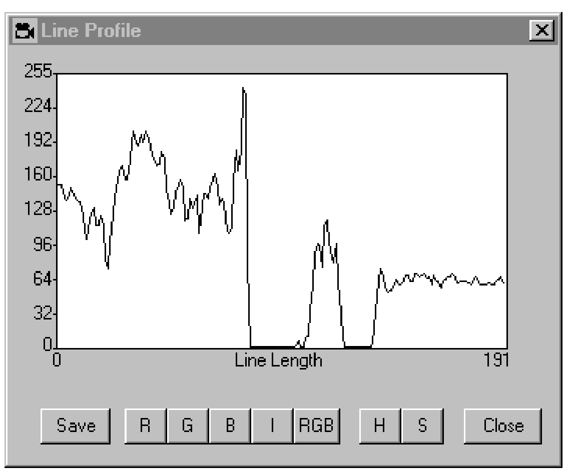

Figure 14. **–** Line Profile dialog box.

R, G, B, I, RGB, H, S buttons – displays a profile of the pixel intensities under the line. The first four buttons stand for red, green, blue, and intensity, respectively. The next button labeled RGB, displays red, green, and blue line profiles all at the same time. The last two buttons display hue and saturation line profiles.

Save – saves the line profile values to a data file.

Delete – deletes the currently highlighted AOI. If only one AOI exists, then that one is deleted.

Delete All – deletes all AOIs.

Next – selecting this menu item (or the corresponding button on the button bar (see figure 2)), will select the next AOI in the sequence of AOIs displayed. When an AOI is selected, it's graphic becomes highlighted and enabled and all other AOIs become disabled. The order of the AOI sequence is the order in which the AOIs were originally created.

Previous – selecting this menu item (or the corresponding button on the button bar (see figure 2)), will select the previous AOI in the sequence of AOIs displayed. When an AOI is selected, it's graphic becomes highlighted and enabled and all other AOIs become disabled. The order of the AOI sequence is the order in which the AOIs were originally created.

AOI Options – the options listed under this menu item are parameters for the particular tracking method selected and, as such, the labels of the menu item change depending on the type of AOI. Thus, if Threshold Tracking AOI is selected, Threshold Tracking parameters will be called from this menu item. Also, any other pertinent options associated with the tracking method would be located here. If multiple AOIs have been selected, then only the options for the one currently highlighted will be displayed. Selection of the AOI Options menu item brings up the appropriate Options dialog box.

Process Sequence – brings up the Image Processing Sequence dialog box (see figure 15) through which a series of image processing operations may be performed inside the selected AOI. The image processing operations are selected and manipulated using the buttons on the right side of the dialog box and are displayed in the list box on the left side of the dialog box. Each of the image processing operations will be executed in the order shown in the dialog box.

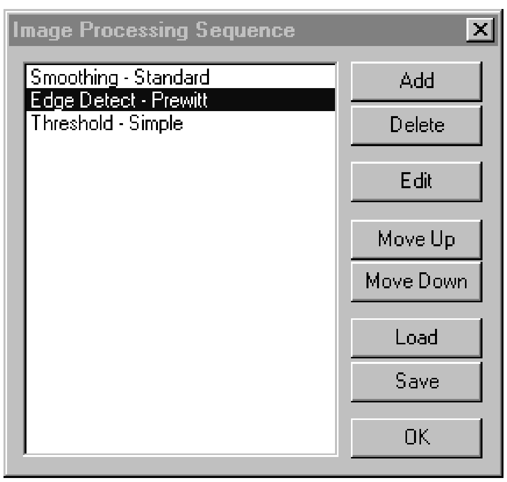

Figure 15. – Image Processing Sequence dialog box.

Add – adds a new image processing operation to current list. Clicking the Add button brings up the Add IP operation dialog box (fig. 16) through which the image processing operation is selected.

Delete – deletes the currently highlighted image processing operation from the list.

Edit – allows the parameters of the selected operation to be changed.

Move Up – changes the order of operations in the list. Clicking this button causes the currently highlighted operation to move up one step.

Move Down – changes the order of operations in the list. Clicking this button causes the currently highlighted operation to move down one step.

Load – loads a previously saved sequence of image processing operations from an ".seq" file. Clicking this button brings up a "Load" dialog box through which the selection is made.

Save – saves the current sequence of image processing operations to a file (".seq"). Clicking this button brings up a "Save As" dialog box through which the selection is make.

The Add IP operation dialog box – this dialog box is displayed when the Add button is selected in the Image Processing Sequence dialog box. Selecting one of these processing operations brings up the pertinent image processing options dialog box described in the Process section below.

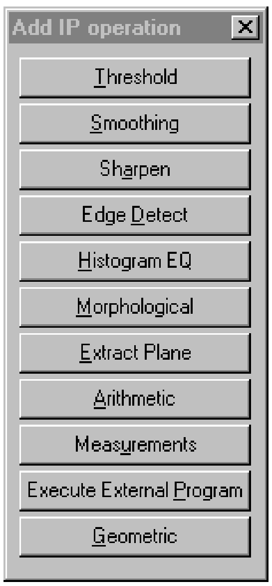

Figure 16. – Add IP operation dialog box.

Update – executes the operations selected in the Process Sequence dialog box and performs the selected tracking operation on the AOI currently selected. This lets the user pre-view the effect of the selections and acts as a check on whether the options and image processing do what you expect them to do.

Process

The menu items under the Process menu perform either image processing operations, measurement operations, or some sort of image transformation. All of the operations operate on an AOI and are designed to perform analysis on a single image frame. The idea behind the Process menu is to perform image processing analysis and testing on one (sometimes two) typical image frame to determine what processing should be done during tracking. Selecting one of the Process menu items will typically bring up a dialog box offering further options.

Threshold – performs an intensity threshold operation. In a threshold operation, the pixels are set either to 0 (or some other intensity) if the original pixel intensity values were below a certain level (threshold), or to 255 (or to some other intensity) if the original pixel intensity values were above a certain level. In the more complex thresholding operations, the pixels are either set to black (0), white (255) or left untouched. Selecting this menu item brings up the Threshold processing dialog box, described below (fig. 17).

Threshold Dialog Box – in the Threshold Dialog Box, the Type parameter selects which threshold filter is used. Depending on the selection, the parameters on the right side of the dialog box, including the two sliders and the three level boxes are enabled or disabled. The value in the left Levels box sets the pixel intensity value below the First Transition. The value in the right Levels box sets the pixel intensity value above the Second Transition (First Transition for Simple Threshold). The value in the middle Levels box sets the pixel intensity value between the two Transitions. The three selections prefaced with the word "Auto" are different from the other threshold selections in that the threshold transitions are automatically calculated. Consequently, for these three selections, the controls on the right side of the dialog box are disabled.

Figure 17. – Threshold dialog box.

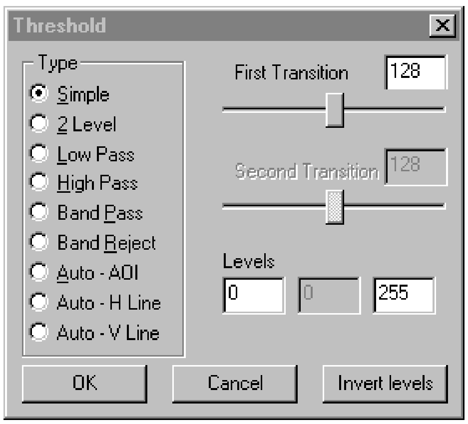

Simple – in the Simple threshold operation (one transition) all pixels in the AOI with intensities below the First Transition are set to the value in the left Levels box, and pixels with intensities above First Transition are set to the value in the right Levels box. The image is then said to be binarized (made up of two intensities).

2 Level – in the 2 level threshold operation (two transitions) the pixels in the AOI with intensities below the First Transition are set to the value in the left Levels box, the pixels between the First and Second Transitions are set to the value in the middle Levels box, and the pixels above the Second Transition are set to the value in the right Levels box.

Low Pass – in the Low Pass threshold operation (one transition), pixels in the AOI with intensities below the First Transition will be "passed" (left unaffected), while pixels above the First Transition will be set to the value in the right Levels box.

High Pass – in the High Pass threshold operation (one transition), pixels in the AOI with intensities above the First Transition will be "passed" (left unaffected), while pixels below the First Transition are set to the value in the left Levels box.

Band Pass – in the Band Pass threshold operation (two transitions), pixels in the AOI with intensities below the First Transition are set to the value in the left Levels box, while pixels above the Second Transition are set to the value in the right Levels box. Pixels with intensities between the First and Second Transitions are "passed" (left unaffected).

Band Reject – in the Band Reject threshold operation (two transitions), pixels in the AOI with intensities below the First Transition and above the Second Transition are “passed” (left unaffected), while the pixels between the two Transitions are “rejected” (set to the value in the middle Levels box).

Auto – AOI – in this operation the threshold transition is automatically calculated. The routine finds the min and max intensities in the AOI and sets the threshold half way. A “Simple” threshold is then performed using the transition just calculated.

Auto – H Line - in this operation the threshold transition is automatically calculated for each horizontal line (row of pixels). For each row, the routine finds the min and max intensities and performs a “Simple” thresholding operation using the value half way between the min and max.

Auto – V Line - in this operation the threshold transition is automatically calculated for each vertical line (column of pixels). For each column, the routine finds the min and max intensities and performs a “Simple” thresholding operation using the value half way between the min and max.

Smoothing – performs a lowpass filtering operation. Smoothing is commonly used for noise reduction. Selecting this menu item brings up the Smoothing Filters dialog box, described below.

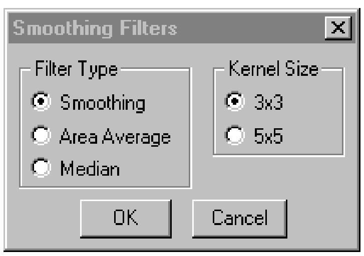

Figure 18. – Smoothing Filters dialog box.

Smoothing – a convolution weighted averaging filter is applied, where each pixel is replaced by a weighted sum of its neighbors. The central pixels are weighted more than the outer pixels. This type of filter is effective in reducing a Gaussian random noise such as the type generated by a digitizer or a tapedeck.

Area Average **–** a convolution non-weighted averaging filter is applied, where each pixel is replaced by a non-weighted sum of its neighbors. All elements of the convolution kernel have the same value (1). This filter can be used as a general blurring (smoothing) filter.

Median **–** the Median filter replaces each pixel with the median in its neighborhood. This type of filter is effective in reducing salt-and-pepper type noise in which a weighted sum filter may create a blotchy effect around each noise pixel.

Kernel Size – the convolution kernel’s size affects the degree of smoothing, with the larger kernel being more severe. This option is not applicable for the Median filter.

Sharpen – performs an edge enhancement filtering operation. Selection of this menu item brings up a Sharpening Filters dialog box, described below.

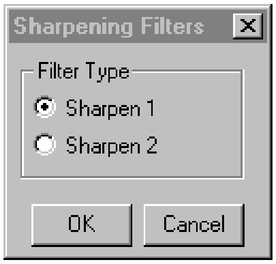

Figure 19. – Sharpening Filters dialog box.

Sharpen 1 – a convolution edge enhancement filter is applied. Between the two sharpening filters available, this one is less severe.

Sharpen 2 – a convolution edge enhancement filter is applied. Between the two sharpening filters available, this one is more severe.

Edge Detect – performs an edge detect filtering operation. Selection of this menu item brings up an Edge Detection Filters dialog box, described below.

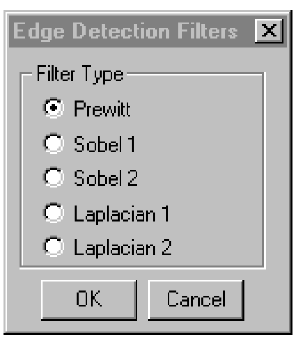

Figure 20. – Edge Detection Filters dialog box.

Prewitt – performs a convolution using a Prewitt edge detect filter. This is the least severe of all edge detect filters, and leaves the thinnest lines.

Sobel 1 – performs a convolution using a Sobel edge detect filter. This method is optimized for speed by using a fast gradient computation approximation as follows: grad = (abs(gradX) + abs(gradY))/2.

Sobel 2 - performs a convolution using a Sobel edge detect filter. This method uses a full gradient computation as follows: grad = sqrt(gradX*gradX + gradY*gradY)

Laplacian 1 – performs a convolution using a Laplacian filter 1.

Laplacian 2 – performs a convolution using a Laplacian filter 2.

Histogram Equalize – performs a histogram equalization (contrast stretching) operation. Histogram equalization alters the pixel values in two ways: one by putting emphasis on pixel intensity values based on the distribution (selected below), and two by shifting intensity values towards the ends of the intensity range, thereby stretching contrast. Selection of this menu item brings up a Histogram Equalization dialog box, described below.

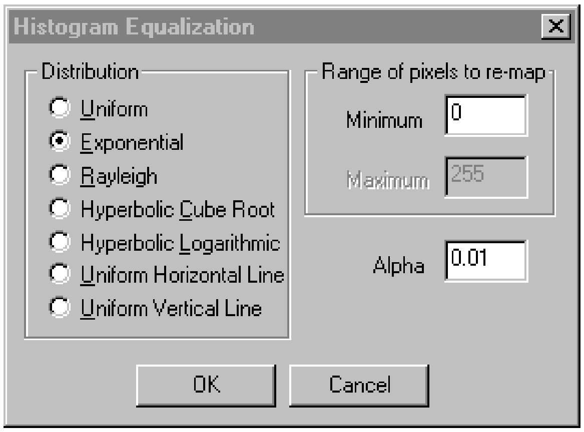

Figure 21. – Histogram Equalization dialog box.

Distribution – refers to the attempted output shape of the histogram curve. Essentially, this means that after the selected equalization filter has been applied, the resulting gray scale histogram distribution will follow (approximate) the given function.

Uniform **–** linearly re-maps the current range of intensities in the AOI to a full intensity range of 0-255.

Exponential **–** re-maps the current range of intensities such that the resultant histogram approximates an exponential distribution, putting emphasis on the bright intensities.

Minimum – specifies the low end of the range of pixels to be remapped. All pixels above this value are remapped into a range between this value and intensity of 255. All pixels below this value are left untouched. The re-mapping is non-linear in which the contrast is stretched more for the brighter pixels and depends on the Alpha parameter.

Alpha – a greater Alpha yields a lower occurrence of the most frequent pixels of the histogram. In other words, a greater Alpha has the effect of shifting emphasis of contrast toward lower intensities, making the image darker. A lower Alpha, since this equalization is exponential, tends to saturate many of the higher intensity pixels.

Rayleigh **–** re-maps the current range of intensities such that the resultant histogram approximates a Rayleigh distribution.

Minimum – specifies the low end of the range of pixels to be remapped. All pixels above this value are remapped into a range between this value and intensity of 255. All pixels below this value are left untouched. The re-mapping is non-linear and depends on the Alpha parameter.

Alpha – a greater Alpha yields a greater occurrence of the most frequent pixels of the histogram. A greater Alpha has (more-less) the opposite effect of the Exponential distribution Alpha, that of shifting emphasis of contrast toward higher intensities, making the image brighter.

Hyperbolic Cube Root – re-maps the current range of intensities such that the resultant histogram is of hyperbolic cube root form. The idea is that the cube root response approximation of a particular detector will make the distribution appear uniform (linear).

Minimum – specifies the low end of the range of pixels to be remapped. All pixels above this value are remapped into a range between this value and the Maximum selected intensity. All pixels below this value are left untouched.

Maximum – specifies the high end of the range of pixels to be remapped. All pixels below this value are remapped into a range between this value and the Minimum selected intensity. All pixels above this value are left untouched.

Hyperbolic Logarithmic – re-maps the current range of intensities such that the resultant histogram is of hyperbolic logarithmic form. The idea is that the logarithmic response approximation of a particular detector will make the distribution appear uniform (linear).

Minimum – specifies the low end of the range of pixels to be remapped. All pixels above this value are remapped into a range between this value and the Maximum selected intensity. All pixels below this value are left untouched.

Maximum – specifies the high end of the range of pixels to be remapped. All pixels below this value are remapped into a range between this value and the Minimum selected intensity. All pixels above this value are left untouched.

Uniform Horizontal Line – re-maps the intensities of each row of pixels based on the Uniform equalization. This routine works well for images where the background intensity varies vertically, since each row is equalized independent of all others.

Uniform Vertical Line – re-maps the intensities of each column of pixels based on the Uniform equalization. This routine works well for images where the background intensity varies horizontally, since each column is equalized independent of all others.

Morphological – performs a morphological filtering operation. Morphological operations are a type of image processing which deal with the structure, connectivity, and geometric relationships between groupings of pixels (objects). They are used for isolating and separating objects from the background so that a particular measurement can be made. The morphological operations supported in Tracker are designed to operate on “binary” images only, so the image must be thresholded before the first morphological operation can be used. Also, Tracker’s morphological operations are based on the convention that black pixels are considered background and white pixels are considered foreground (the object). Selecting this menu item brings up the Morphological dialog box, described below.

Figure 22. – Morphological dialog box.

Erosion – in this operation, any pixel whose neighborhood is not completely white, is changed to black. This operation has the effect of changing white edge pixels to black, thereby decreasing the size of an object. If two or more objects are connected, erosion has the effect of separating the objects.

Dilation – in this operation, any black pixel that has at least one white neighbor is changed to white. This operation has the effect of changing black edge pixels to white, thereby increasing the size of an object. If two or more objects are separated by a gap, dilation has the effect of fusing the objects together.

Thinning – erodes an object making sure not to disconnect any currently connected groups of pixels. However, only specific pixels are eroded away. The sides of an object are eroded in the direction of a line equidistant from its nearest outer boundaries. Given enough iterations, this operation will erode to a skeleton of the original object.

Thickening – dilates an object making sure not to fuse any currently disconnected groups of pixels.

Opening – a shortcut for erosion followed by dilation. This combination of operations is so common that it has been given a special name; opening. It is useful for cleaning up (eroding to nothing) an image by eliminating small particles but leaving the large ones. For multiple iterations, this function is equivalent to first performing multiple iterations of erosion followed by multiple iterations of dilation.

Closing – a shortcut for dilation followed by erosion. This combination of operations is so common that it has been given a special name; closing. It is useful for fusing small particles together into large ones, such as connecting parts of a broken object. For multiple iterations, this function is equivalent to first performing multiple iterations of dilation followed by multiple iterations of erosion.

Hole Fill – fills in any holes in an enclosed grouping of pixels (object).

Border Kill – eliminates (changes to black) all pixels that are connected to pixels touching the border of an AOI. Any object separated by a gap from an AOI border will be left untouched.

Hole Extract – swaps the pixel value between a hole and the object. Thus, pixels what were white (object) are changed to black and pixels that were labeled as part of a hole (black) are changed to white.

Reconstruct – attempts to "reconstruct" an object from a given set of markers. These markers are the results or leftovers of an eroded image. This algorithm works in several steps. First, several iterations of erosion are applied to the image until the small, unwanted objects (crud) are eroded away. What is left of the erosion is called a marker image and the objects in the marker image are called markers. Second, the marker image is forwarded to the reconstruct algorithm. The original image is then superimposed over the marker image. And third, all objects touching any of the markers are kept, and all others are eliminated.

Skeleton – performs a thinning operation until a skeleton of the object is obtained. A skeleton is a result of a thinning operation until the object is reduced to single pixel thick lines.

Outline – generates an outline of an object. All interior pixels of an object are set to black and only the edge pixels are left untouched, thereby creating an outline.

Iterations – selects the number of times the particular operations is to be performed repeatedly. It is disabled when not applicable.

Extract Plane – extracts a color plane from a 24-bit image. It should be noted that extracted plane (which is 8-bit) is then converted to 24-bits so that it can be reinserted back into the original image for viewing, however it will look and act as grayscale (8-bit) to any further image processing operations. Selecting this menu item brings up the Extract Color Plane dialog box, described below.

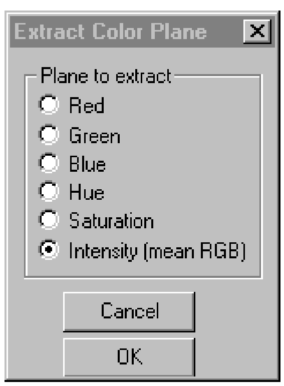

Figure 23. – Extract Color Plane dialog box.

Red **–** extracts a red (8-bit) plane from a color (24-bit) image.

Green – extracts a green (8-bit) plane from a color (24-bit) image.

Blue – extracts a blue (8-bit) plane from a color (24-bit) image.

Hue – extracts a hue (8-bit) plane from a color (24-bit) image. *

Saturation – extracts a saturation (8-bit) plane from a color (24-bit) image. *

Intensity – extracts an intensity 8-bit plane from a color (24-bit) image. The intensity is obtained by a simple average of the three color planes (R+G+B)/3. *

* Based on RGB to HSI conversion algorithms by Strickland, Kim, and McDonnell, "Luminance, hue, and saturation processing of digital color images" in *Applications of Digital Image Processing IX*, A. G. Tescher, ed., Proc. Soc. Photo-Opt. Instrum. Eng. 697, 286-292 (1986)

Arithmetic – performs an arithmetic operation between two images or between an image and a constant. If the operation is between the current image and a constant, every pixel of the image is evaluated relative to the constant. If the operation is between the current image and a second image (obtained from a disk file), every pixel of the current image is evaluated relative to the corresponding pixel in the second image. Both images must be the same size. Some of the operations, such as the logical operations, are more often used with binary (thresholded) images, while others are better used with grayscale. Selecting this menu item brings up the Arithmetic dialog box, described below.

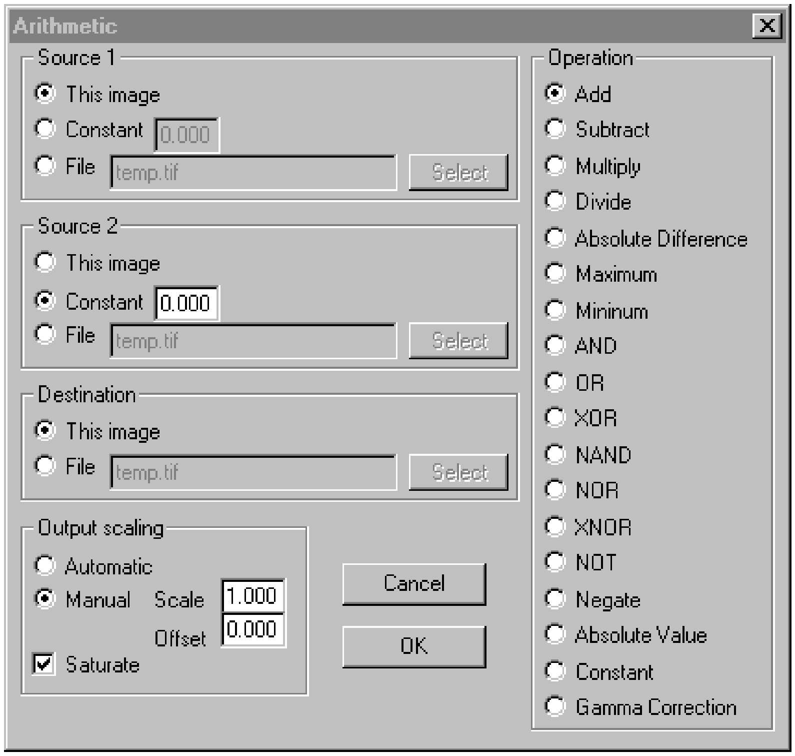

Figure 24. – Arithmetic dialog box.

Source 1 – selects the first of the two inputs to the operation. There are three possible selections for source 1. “This image” selects the image currently contained within the AOI. ”Constant” selects the input to source 1 to be a constant, which is typed in the box on the right. “File” selects an external file from disk (TIF only).

Source 2 – selects the second of the two inputs to the operation. There are three possible selections for source 2. “This image” selects the image currently contained within the AOI. ”Constant” selects the input to source 2 to be a constant, which is typed in the box on the right. “File” selects an external file from disk (TIF only).

Destination – selects where to put the result of the computation. Typically it is displayed back to the current image, but a second option is to save it to a file, possibly to be used by another operation later.

Output scaling – selects type of scaling to be performed on the resultant image. This is because many of the operations can produce a result outside of the normal intensity range of 0-255. The “Automatic” scaling performs a uniform histogram equalization on the result, which scales the intensities. The “Manual” scaling modifies the result based on the “Scale” and “Offset” options. The Scale simply multiplies the result by the value in the box and the Offset adds or subtracts by the value in its box. The “Saturate” option clips the resultant intensity at 0 or 255.

Operation **–** these options select which Arithmetic operation to perform between source 1 and source 2. The operation is performed between every pixel of source 1 and source 2.

Add – adds source 2 to source 1.

Subtract **–** subtracts source 2 from source 1.

Multiply – multiplies source 1 by source 2.

Divide – divides source 1 by source 2.

Absolute Difference **–** subtracts source 2 from source1 and returns the absolute value of the subtraction.

Maximum – returns the maximum value between source 1 and source 2.

Minimum – returns the minimum value between source 1 and source 2.

AND **–** performs a logical AND operation between source 1 and source 2. For two binary images, this operation returns the intersection of the two images (overlap between the objects from both images).

OR **–** performs a logical OR operation between source 1 and source 2. For two binary images, this operation returns the union (objects in both images as well as the overlap) of the two images.

XOR **–** performs a logical XOR operation between source 1 and source 2. For two binary images, this operation returns the exclusive OR (objects in both images but not the overlap) of the two images.

NAND – performs a logical NAND operation between source 1 and source2. For two binary images, this operation is equivalent to an AND operation followed by a NOT operation.

NOR **–** performs a logical NOR operation between source 1 and source 2. For two binary images, this operation is equivalent to an OR operation followed by a NOT operation.

XNOR – performs a logical XNOR operation between source 1 and source 2. For two binary images, this operation is equivalent to a XOR operation followed by a NOT operation.

NOT – performs a logical NOT operation on source 1. A NOT operation basically flips the image intensities about the value 127.5, thus inverting the intensities.

Negate **–** performs a negate operation on source 1. This operation produces results similar to the NOT function, except for binary images, but arrives at the result in a different way; by subtracting a value of 256 from the current intensity and then taking an absolute value.

Absolute Value – takes an absolute value of source 1. Usually it is used in conjunction with a previous operation, which might produce a negative result.

Constant **–** sets the entire AOI region to the selected constant intensity value.

Gamma Correction **–** raises each pixel to the power ***gamma***. It is used to correct for errors in the gamma settings of cameras or displays. The value (gamma) is entered into the constant box of source 2 and the value of gamma typically is a small number close to 1.

Execute External Program – this menu item allows Tracker to perform several useful operations such as saving AOI(s) to disk, loading AOI(s) from disk, and executing an ancillary program outside of Tracker to process an AOI. Selecting this menu item brings up the External Program Interface dialog box, described below.

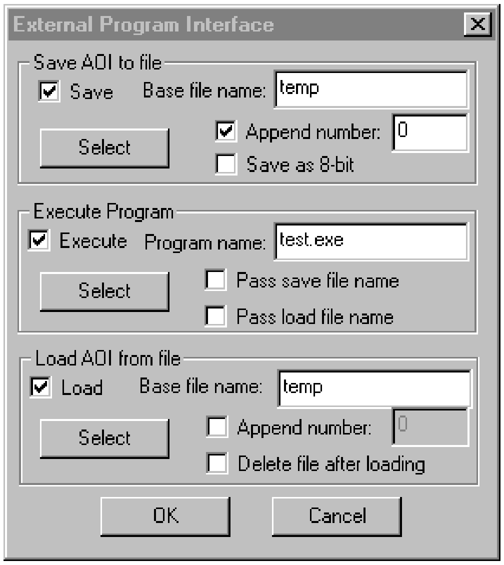

Figure 25. – External Program Interface dialog box.

Save AOI to file – allows Tracker to save AOI(s) to a TIF file on disk. With the "Append number" option selected, Tracker can save a numbered sequence of TIF files. This option can also be used for temporarily saving an intermediate result to a file so that it can be picked up later in the program for further processing.

> ***Save*** – enables the capability to save an AOI to a file.
>
> **Select** – clicking this button brings up the Save As dialog box through which a base file name and path can be selected. Note that if a filename with an extension of ".tif" is selected, the extension will be stripped off since base file name does not contain an extension. A ".tif" extension will be added automatically by the routine. The selected file name appears in the "Base file name" box.
>
> ***Base file name*** – a name of a file (and path) to which the AOI is saved. The term "Base file name" refers to the part of the filename ***without*** the appended numbers and the file extension. The numbers are appended if the "Append number" option is used and the file extension is appended automatically. The name (and path) can be either typed in by hand or selected with the "Select" button.
>
> ***Append number*** – enables numbering to be appended sequentially to the base file name, starting with the number typed here.
>
> ***Save as 8-bit*** – specifies the AOI to be saved as an 8-bit (grayscale) image even if the current image is color. This option is intended as a disk space saving feature since Tracker digitizes even monochrome video to 24-bits.

Execute Program – allows Tracker to execute an external program (from within Tracker) and pass results (image) to it, or obtain results from it. This capability might be useful when some computation is required that is not supported by Tracker.

> ***Execute*** – enables the capability to execute an external program.
>
> ***Select*** – brings up the Load dialog box through which a program name and path can be selected. The selected program name appears in the "Program name" box.

Program name – the name (and path) of the program to be executed. The name can be either typed in by hand or selected with the "Select" button. If an image input to the program is required, it can be specified by explicitly typing the image name after the program name or it can be automatically specified using the "Pass save file name" and "Pass load file name" options.

Pass save file name – enables an automatic appending of image file name from the "Save AOI to file" section of this dialog box to the program name. For example, if the Save base file name (as shown in the dialog box) is "temp" and the "Append number" is 0, then the name temp0.tif will be appended to the Program's command line – i.e., "test.exe temp0.tif". By having the "Append number" selected, a new AOI (temp0.tif, temp1.tif, etc.) will be passed to the program each time this function is executed.

Pass load file name **–** enables an automatic appending of image file name from the "Load AOI from file" section of this dialog box to the program name. For example, if the Load base file name (as shown in the dialog box) is "temp" and no "Append number" is selected, then the name temp.tif will be appended to the Program's command line – i.e., "test.exe temp.tif".

Load AOI from file – allows Tracker to load an image from a TIF file on the hard disk. With the "Append number" option selected, Tracker can load a numbered sequence of TIF files from disk. Each time this function (Execute External Program) is executed, a new image is loaded. One requirement is that the image file loaded must be the same size as the current AOI (the image in the AOI will be replaced by the image in the file).

Load – enables the capability to load an AOI from a file.

Select **–** brings up the Load dialog box through which a base file name and path can be selected. Note that if a filename with an extension of ".tif" is selected, the extension will be stripped off since base file name does not contain an extension. A ".tif" extension is be added automatically by the routine. The selected file name appears in the "Base file name" box.

Base file name – a name of a file (and path) to be loaded. The term "Base file name" refers to the part of the filename ***without*** the appended numbers and the file extension. The numbers are appended if the "Append number" option is used and the file extension is appended automatically. The name can be either typed in by hand or selected with the "Select" button.

Append number - enables numbering to be appended sequentially to the base file name, starting with the number typed here.

Delete file after loading – specifies that the file will be deleted after it has been loaded. This option is designed for cleaning up after a temporary file was saved and is no longer needed.

Measure – performs area, centroid, and pixel intensity measurements. For area and centroid measurement, if a number of separate objects exist inside an AOI, the user has the option to choose a total area for all of the objects, separate areas of each of the objects, or the area of just the largest one. The pixel intensity measurement measures max, min, and average of pixels inside the AOI. The area and centroid measurements are based on a binary (thresholded) image, therefore for these two operations, thresholding must be done before the measurement operations are performed. However, thresholding should not be performed before the Max-Min-Average operation because this measurement measures grayscale information. Selecting this menu item brings up the Measurement dialog box, described below.

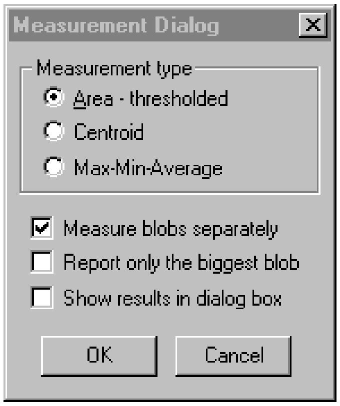

Figure 26. – Measurement dialog box.

Area – thresholded – measures an area of an object(s) based on a number of non-zero pixels. Since the routine is designed for binary operation, the AOI should be thresholded first. The routine simply counts up the number of thresholded pixels. By default, Tracker reports a number of blobs (objects) in the AOI and total area. Area measurement of individual blobs can be selected using the “Measure blobs separately” option.

Centroid – measures centroid locations of blobs (objects) in an AOI. Since the routine is designed for binary operation, the AOI should be thresholded first. By default, Tracker reports a number of blobs (objects) in the AOI and centroid for all of the blobs. Centroid measurement of individual blobs can be selected using the “Measure blobs separately” option.

Max-Min-Average – measures the maximum intensity, minimum intensity, and average intensity of pixels in the AOI.

Measure blobs separately – enables blob measurement for individual blobs rather than a cumulative measurement for all the blobs in the AOI. If selected, blobs will be measured individually. If not selected, a cumulative (total) measurement for the entire AOI is made. This option works in conjunction with the check box labeled “Report only the biggest blob”, which is described below.

Report only the biggest blob – enabled only if the check box above (“Measure blobs separately”) is checked. If this option is selected, measurement for only the biggest blob is performed. If not selected, individual measurement of every blob will be performed.

Show results in dialog box – the result of the measurement is automatically written to the current results file. Selection of this option displays the results in a dialog box, in addition to the result file.

Geometric – currently the only geometric function supported is rotation. Selecting this menu item brings up the Geometric Functions dialog box, described below

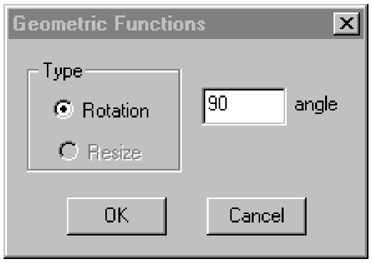

Figure 27. – Geometric Functions dialog box.

Rotation – rotates the area inside the current AOI by the angle specified. Note that part of the area will be cropped if the AOI is not square.

Undo – the Undo menu item allows the user to "undo" any image processing operation performed under the Process menu. If multiple imaging operations were performed, then multiple execution of the Undo menu item will undo the imaging operations one by one. This function has a corresponding button on the button bar (see figure 2).

Redo – the Redo menu item allows the user the "redo" any operation that was undone with the Undo operation. In this way the Redo and Undo are opposites of one another. Also, if an image processing operation was performed and not Undone, the Redo operation will execute another iteration of the imaging operation. This function has a corresponding button on the button bar (see figure 2).

Tracking

The menu items under the Tracking menu deal with performing tracking and managing how data is stored in an output file.

Track 1 Frame – selecting this menu item or clicking the corresponding button on the button bar causes Tracker to track one frame and then stop. This function is sometimes useful to check if everything is proceeding as expected. Once the user is convinced that everything is working well, then continuous tracking may be employed.

Track Continuously – selecting this menu item or clicking the corresponding button on the button bar causes Tracker to start the continuous tracking process. This menu item works like a toggle, thus a second execution of this menu item (or clicking the button a second time) causes tracking to stop.

Result File – brings up a Results File dialog box (see figure 28) through which an output data file can be selected (into which the tracking results are written). If an output data file is not explicitly selected (through this dialog box), Tracker will write the tracking data to a default file named "results.txt" in the directory from which Tracker is run. The data is written to a file in one of two possible formats, depending on the type of information saved. Tracker does not save the actual data to a file until the tracking process has been stopped. During tracking, the tracking results are written to RAM and only when tracking is terminated, does Tracker make a decision on how it should be formatted. This allows Tracker to make a decision on the format after all of the data has been collected and allows presenting the data in the most legible format. The first format is the "column format" in which all data collected for each frame fits on the same row. In this format, all data is formatted into columns of numbers, one row per time frame. The second format is the "wide format" and it is used for complex and varying types of data, which would not fit neatly into columns (such as output of snake tracking when positions along the line are requested).

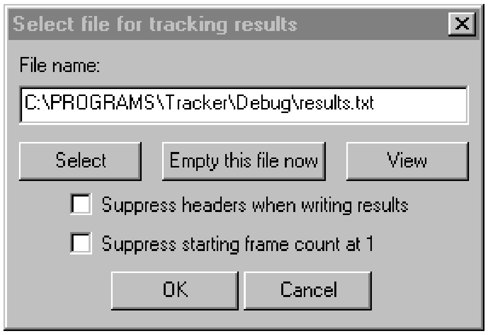

Figure 28. – Tracking results selection dialog box.

File name – name and path of the result (output data) file. It can be typed by hand or it can be selected with the "Select" button.

Select – brings up a Save As dialog box through which the result file name and path can be selected. The path and file name selection is displayed in the "File name" box above.

Empty this file now – deletes the contents of the currently selected result file. The user is prompted one more time before the action is executed.

View – displays the contents of the result file in Windows Notepad.

Suppress headers when writing results – selects whether or not to write header information each time tracking is started. Whenever tracking is initiated, by default, a new header is written to the output data file. The header consists of labels of what each data means. Selecting this option will stop writing the headers to the file, writing only the data. This feature is useful when tracking is started and stopped several times and additional headers after the first one are not desired.

Suppress starting frame count at 1 – selects whether or not to start the frame count at 1 each time tracking is started. Whenever tracking is initiated, by default, the frame number starts at 1. Selecting this option will continue the frame count where it was last stopped. This feature is useful when tracking is started and stopped several times and it is desired that the frame number start where it left off.

Help

The menu items under the Help menu provide access to the Tracker documentation and version information.

View Documentation - This menu opens the Tracker documentation file (Tracker.pdf) with the default program your computer uses to view Adobe PDF files. If you have no program to view PDF files, you may download the free Adobe Acrobat Reader at the following URL: http://www.adobe.com.

About - The menu opens the "About box", which shows copyright and author information, and most importantly, the date that your version of Tracker was compiled. New minor versions of Tracker 3 are compiled as bugs are fixed.

Appendix A

Digit - The neural network weight file generator

Introduction

The neural network character recognition method used in Tracker reduces each sub-AOIs containing a character (representing numbers only) to a low resolution image (7 rows by 5 columns, for example), then presents this image to a neural network for evaluation to find the closest match for the character. The neural network uses a "weight file" to evaluate the image. The weight file is the result of training the network on example images. Training the network can take a long time, but evaluating a trained network is very fast. For this reason, tracker only evaluates the neural network; and a separate program does the training before Tracker is run. That program is "Digit".

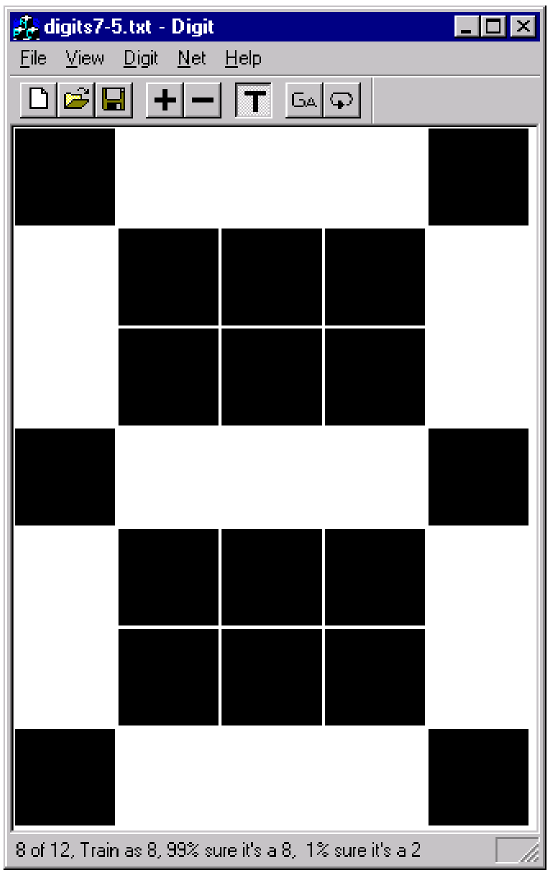

Figure 29. Digit program used for training a neural net to recognize a set of numbers.

Tracker comes with two pre-generated weight files: "seven.wt" (a weight file for a 1 row by 7 column network, trained to identify 7 segment LED numeric displays), and "digits7-5.wt (a 7 row by 5 column network, trained to identify numbers). If you wish to recognize numbers in a strange font, or use reduced resolutions other than 7 by 5, then you will need to use Digit to generate the weight file trained on your new data set.

Creating a Training Set

When you first start the digit program, you will see a grid of black cells (it is helpful to resize the digit window so that the cells are approximately square). The default grid size is 7 rows by 5 columns, so you will see 35 black squares arranged in a 7 by 5 grid. There is a menu bar and a button bar at the top of the window, a cell display area in the middle of the window, and a status bar at the bottom of the window.

The size of the grid (the number of rows and columns) can be changed with the "Digit→Set Rows and Columns" menu items.

The cells represent pixels in the reduced resolution image of the character. You can edit the image by clicking the mouse on a cell; this toggles the cell between black and white. Black is the background color, and white represents a pixel that is "on" for a particular character. The diagram shown above shows the character "8" as it might appear in a 7 by 5 grid.

After the image is edited to show the image of the character you want the neural network to evaluate, you must tell the network what number this character is supposed to represent. This is done with the "Digit→In Training Set" menu item (or the equivalent button on the button bar - the one with the big "T"). If an image is not added to the training set in this manner, the image will be evaluated by the neural network (useful for testing), but it will not be used for training.

An image must be created for each character in the set to be recognized. New images are created with the "Digit→New Image" or the "Digit→Duplicate Image" menu items. Images are delete with the "Digit→Delete Image" menu item. The "Digit→Next Image" and "Digit→Previous Image" menu items (or the equivalent button with the "+" and "-" symbols on them) cycle through all of the images that have been created.

The "File→Open" and "File→Save" menu items (or the equivalent buttons) are used to open and save the set of images on disk.

The status bar at the bottom of the digit window shows how many images are in the image set, what value will be used for training the current image (if any), and the top two candidates numbers that the current image would evaluate to using the current network weights. Initially the weights are random, so the image is equally likely to evaluate to any number.

Training the Network

Training the network is started (and stopped) with the "Net→Train Continuously" menu item (or the equivalent button that looks like an arrow bent back upon itself). While training the status bar shows information about how well the training is progressing. The important number on the status bar is the RMS error, which should approach zero when the net is fully trained.

At any time, training may be stopped and the progress examined. Cycling through each image will show a training summary in the status bar (something like "97% sure it's a 2"). The "Net→Evaluate All" menu item can be used to evaluate all of the training images and display a summary of average sureness and the number of misclassified images.

A large set of training images can sometimes cause the standard back-propagation algorithm that the neural network uses for training to get stuck in a local minimum for long periods of time. If the RMS error does not seem to be decreasing, the network can switch to genetic algorithm based training. This is toggled with the "Net→Genetic Algorithm Training" menu item (or the equivalent button that says "Ga"). If training fails for some reason, the network can be re-initialized to a random state by the "Net→Randomize Weights" menu item. Training can then be attempted again.

The weights that define the trained network are saved to a file that Tracker can use with the "Net→Save Weights" menu item. A set of saved weights can also be reloaded to evaluate a set of images (or for further training) with the "Net→Open Weights" menu item.

Appendix B

Hardware notes

The object tracking system consists of a number of discrete components as illustrated in figure 30. The arrows on the lines connecting the components show the image signal flow. Each of the image devices are also connected (not shown in fig. 30) to the PC using serial (RS-232) interface. A "smart" serial card enables multiple ports to share the same interrupt thus providing separate serial lines.

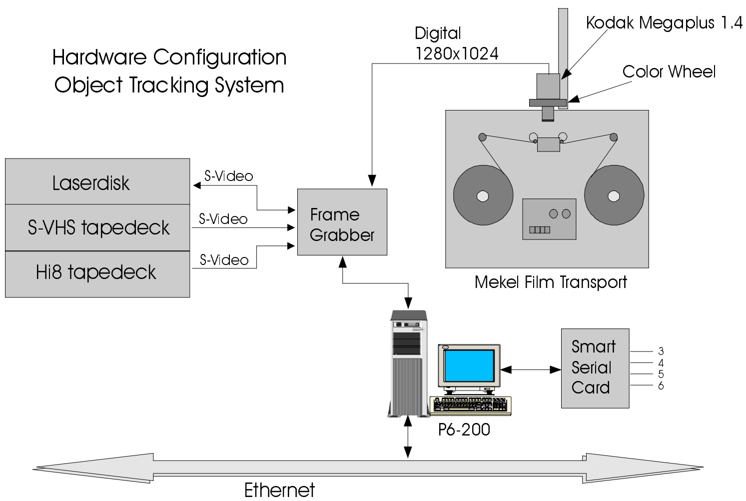

Figure 30. Tracking system hardware layout.

Disk files

Tracker can accept input from digital file saved on disk. The file formats that Tracker can read are as follows: tif, gif, tga, dxf, eps, img, jpg, pcx, png, wmf, wpg, PhotoCD, and bmp. Currently, Tracker can save images only in a tif file format.

Matrox Meteor frame grabber

The Matrox Meteor/PPB frame grabber is a low cost PCI-based frame grabber, which is capable of acquiring color images in Y/C (S-Video), or NTSC format. The images are digitized to a resolution of 640x480 pixels. Note that the Meteor frame grabber is used with Pentium computers while the Meteor/PPB frame grabber is used with Pentium Pro and Pentium II computers.

Mekel film transport

The Mekel film transport system consists of a backplate which houses the motors and electronics and the film transport head which performs the actual film movement. Two head sizes were purchased supporting both 16mm as well as 35mm film. Both heads are pin-registered thus preventing any film slippage or rotation in the gate. The heads, due to their design, have a convenient feature of allowing the camera to look down directly at the film without having to resort to mirrors or prisms. A detailed manual is included with the system; therefore only the necessary steps to get a user started are included here.

A Kodak Megaplus 1.4 camera is used to image the film. The Megaplus 1.4 is a monochrome digital camera with a resolution of 1335x1024. Only 1280x1024 are usable, as that is the maximum size handled by the frame grabber. The frame grabber has a special high-speed digital connector that the camera uses. The maximum frame rate of the camera is 6 f/s, but that speed is rarely realized in practice as it is slowed down by the shutter speed. The camera head is attached to a 3-axis translation stage and the camera control unit (CCU) is located below the film transport. The CCU is controlled by the computer through the frame grabber. Up to a 2:1 zoom can be accomplished by moving the camera head up or down using the large screw knob at the top of the translation stage. *Note: the Megaplus' ON/OFF switch is located in the back of CCU on the left-hand side.*

The black box between the Kodak Megaplus head and the lens is a colorwheel. It was custom made by Forster Engineering. The colorwheel is used to generate a color image by rotating color filters (red, green, blue, or clear) in front of the camera before each frame is acquired. The colorwheel can be controlled from Tracker using a serial interface. Although it can also be operated manually by pressing the small button on the front face, this method should NOT be used as it is meant for testing purposes only. Since most of the lamp energy is weighted towards the red (longer wavelengths) the computer compensates (balances colors) by adjusting the exposure duration for the four planes.

A macro lens, designed to work with a Mekel film transport system, is positioned at the bottom of the colorwheel housing. The extra large focal range enables it to focus throughout the length of travel when zoom is employed. *Note: a lens extender ring, between the colorwheel and the lens, must be inserted when 16mm transport head is used and taken out when 35mm transport head is used.*

Laserdisk (Sony LVR-3000AN)

The laserdisk is a real-time WORM (Write Once Read Many) type video recorder device. It uses cartridges with a capacity of 24-minutes of live video per side. The laserdisk responds much faster to computer issued commands than a tapedecks and therefore automated tracking can proceed at a much faster rate. Other nice features include output of full frame (or a field) when single stepping through frames, random access to image frames, and a more permanent storage than videotape (and thus can be used for video archival.

Re-writable Laserdisk (Panasonic LQ-4000)

The LQ-4000 is a laserdisk unit, which uses rewritable cartridges capable of recording 30 minutes of real-time video. The LQ-4000 is similar to the Sony LVR-3000N in most of its capabilities including outputting a full frame, random access, and quick incrementation of frames.

Hi8 Video tape (Sony EVO 9650)

Although this tapedeck can be can be controlled through the RS-232 interface, it does a rather poor job handshaking. Also the image seems to shake for about a second and a half after a frame is advanced. For these reasons a 2-second delay is used between frames (unfortunately slowing down the tracking operation).

VHS & S-VHS Video player (JVC 525)

The JVC-525 tapedeck is the only tapedeck on the market, at this time, which employs dynamic tracking (utilizes a special head to keep track of frame positions) ensuring the frames to be aligned on the screen without the interference of a scroll bar. The built-in time-base corrector ensures that the frames are stable. It is set to auto-detect time code and uses it if present. As with any tapedeck, when in Pause mode as is the case when tracking, only a field is displayed, not a frame.

VHS & S-VHS Video tape (Panasonic 7300)

This tapedeck may be used when the JVC 525 is not available. The 7300 is a very poor device from which to track and thus it is encouraged that the image data is copied to laserdisk and tracking is done there.

Smart Serial Card - Rocket Port 4

The Rocket Port is a four-port serial card, which manages all four ports with a single interrupt. This expands the number of available serial ports from 2 to 6.

Appendix C

Character Recognition AOI

There are two separate character recognition AOI issues that need further explanation. One is the positioning of the AOI using a mouse. The character recognition AOI is quite complex, employing control over resizing, rotation, skewing, and gap size. Second, the two character recognition algorithms implemented by Tracker require different AOI placement over the character string. Both issues are explained below. Also, as mentioned in the Reference Guide, the character recognition is currently set up only to recognize numbers (digits) from 0 to 9.

AOI position controls

The character recognition AOI is controlled using the mouse – although some of the parameters can also be specified through the character recognition dialog box. Figure 31 shows the functionality of different control points (AOI corners) on a character recognition AOI that contains four sub-AOIs. The AOI is manipulated such that the character string fits inside the AOI, one sub-AOI per each character.

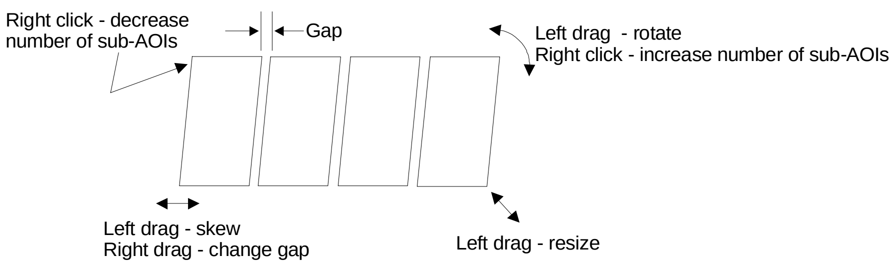

Figure 31. – Character recognition AOI controls.

AOI placement vs. recognition method

Both of the character recognition methods – MIL OCR and Neural Network – have definite requirements on placement of the AOI so that the character string can be properly recognized.

MIL OCR Method

The MIL OCR method requires the AOI to be placed over the character string such that some space is left around each of the characters. Artifacts in the image, such as the period after the number 2, have little effect on the recognition accuracy. Thus, the sub-AOI gap width is not important. See figure 32 for proper AOI placement.

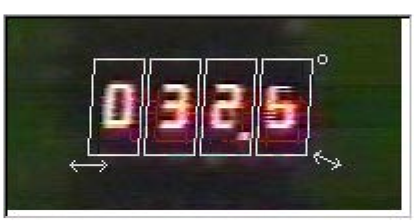

Figure 32. – AOI placement for the Mil OCR method.

Neural Network Method

The Neural Network method requires the AOI to be placed such that the left edge of each sub-AOI is up against the edge of the character. The other three sides do not matter whether there is space around them or not. Since this method is affected by artifacts such as a period after a character, changing the gap width can be used to eliminate the artifact from the sub-AOI.

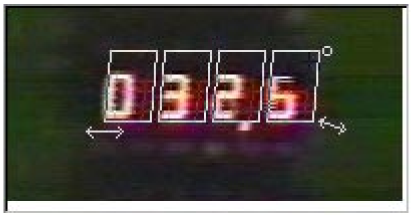

Figure 33. – AOI placement for the Neural Network method.